
PHD THESIS

UNDERSTANDING THE BEHAVIOR OF BELIEF PROPAGATION

Convergence Properties, Approximation Quality, and Solution
Space Analysis

conducted at the
Signal Processing and Speech Communications Laboratory
Graz University of Technology, Austria

by
Dipl.-Ing. Christian Knoll, BSc.

Supervisors:
Assoc.Prof. Dipl.-Ing. Dr.mont. Franz Pernkopf

Assessors/Examiners:
Assoc.Prof. Dipl.-Ing. Dr.mont. Franz Pernkopf
Dr. Adrian Weller

Graz, November 8, 2019

Statutory Declaration

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly marked all material which has been quoted either literally or by content from the used sources. The text document uploaded to TUGRAZonline is identical to the present doctoral dissertation.

date

(signature)

Contents

Abstract	9
Acknowledgments	11
Notational Conventions	13
1 Introduction	19
1.1 Motivation	19
1.2 Five Relevant PhD Theses	21
1.3 Contribution and Outline	22
2 Background	25
2.1 Probability Theory	25
2.2 Graph Theory	26
2.3 Probabilistic Graphical Models	27
2.3.1 Undirected Graphical Models (Markov Random Fields)	28
2.4 Binary Pairwise Models	29
2.4.1 Model Description	29
2.4.2 Generality of Binary Pairwise Models	31
2.4.3 Exponential Representation and Model Parametrization	32
2.4.4 The Ising Model	33
2.5 Inference	34
2.5.1 Exact Inference: Efficient Methods	36
3 Approximate Inference: Belief Propagation	41
3.1 Motivation	41
3.2 Preliminaries	42
3.3 Variational Interpretation of BP – A Physics Perspective	45
3.3.1 Energy Landscape of the Bethe Free Energy	49
3.3.2 The Bethe Approximation and Belief Propagation	49
3.3.3 Other Variational Approaches	52
3.4 Approximation Quality	53
3.5 BP as a Dynamical System	55
3.5.1 Enhancing Belief Propagation’s Properties	56
3.5.2 Initialization	56
3.5.3 Update	57
3.5.4 Evaluation	60
3.6 Conclusion	62
4 Discrete-Time Dynamical Systems: Solution Space Analysis	63
4.1 Discrete Time Map	63
4.2 Finding Fixed Points of a Discrete Time Map	64
4.2.1 Systems of Polynomial Equations	64
4.3 Solving Polynomial Systems: A Comparison of Methods	65
4.3.1 Numerical Methods	66
4.3.2 Symbolic Methods	66
4.3.3 Numerical Polynomial Homotopy Continuation (NPHC) Method	70
4.4 Stability Analysis of Fixed Points	73
4.4.1 Stability Analysis for a Nonlinear Discrete Time Map	74

5	Solution Space of Belief Propagation: Number of Fixed Points and Their Stability	77
5.1	Motivation	77
5.2	Fixed Points of Belief Propagation	79
5.2.1	Number of Fixed Points on Regular Ising Graphs	79
5.3	Fixed Point Equations of Belief Propagation	80
5.3.1	Solving the Fixed Point Equations	81
5.3.2	Polyhedral Homotopy Method	82
5.4	Stability of Fixed Points	85
5.4.1	Reformulation of Belief Propagation for Binary Variables	85
5.4.2	Linearization	86
5.5	Selected Models with Ising Potentials	87
5.6	Number of Fixed Points and Marginal Accuracy	88
5.6.1	Evaluation Criteria	88
5.6.2	Grid Graphs with Random Factors (Spin Glasses)	89
5.6.3	Grid Graphs with Uniform Factors	90
5.6.4	Complete Graphs with Uniform Factors	91
5.6.5	Fixed Point Evolution	92
5.6.6	Runtime Analysis	94
5.7	Empirical Stability Analysis of Belief Propagation on Ising Models	95
5.7.1	Vanishing Local Field	95
5.7.2	Non-Vanishing Local Field	95
5.8	Theoretical Stability Analysis	98
5.9	Application: Error-Correcting Codes	101
5.10	Conclusion	103
6	Enhancing Belief Propagation: Self-Guided Belief Propagation	105
6.1	Motivation	105
6.2	Self-Guided Belief Propagation (SBP)	106
6.2.1	Practical Considerations	107
6.2.2	Pseudocode	107
6.3	Experiments	108
6.3.1	Experimental Settings	109
6.3.2	Attractive Models	110
6.3.3	General Models	111
6.4	Theoretical Analysis of Self-Guided Belief Propagation	112
6.4.1	Definitions	112
6.4.2	Properties of Self-Guided Belief Propagation	113
6.5	Conclusion	116
7	Understanding Belief Propagation: Accurate Marginals or Partition Function	119
7.1	Motivation	119
7.2	Model Specifications	120
7.2.1	Definitions	120
7.2.2	Patch Potential Models	121
7.3	Fixed Point Behavior	122
7.3.1	Combination Of Fixed Points	122
7.3.2	Solution Space of Patch Potential Models	123
7.3.3	Approximation Accuracy	124
7.3.4	Comparison Of Marginal Accuracy and Partition Function Accuracy	125
7.4	Theoretical Analysis	126
7.4.1	Effective Field	127
7.4.2	Definition of Region (<i>II</i>)	127

7.4.3	Properties Of Region (<i>II</i>)	127
7.5	Conclusion	133
8	Discussion and Open Questions	135
A	Appendix	139
A.1	Proofs from Chapter 6	139
A.1.1	Proof of Lemma 4	139
A.1.2	Proof of Theorem 9	140
A.2	Proofs from Chapter 7	143
A.2.1	Proof of Theorem 11	143
A.2.2	Proof of Corollary 15.1	144
A.2.3	Proof of Theorem 16	145
A.3	Pseudocode	146
B	Guide to Quotes	147

Abstract

Probabilistic graphical models are a powerful concept for modeling high-dimensional distributions. Besides modeling distributions, probabilistic graphical models also provide an elegant framework for performing statistical inference; because of the high-dimensional nature, however, one must often use approximate methods for this purpose.

Belief propagation performs approximate inference, is efficient, and looks back on a long success-story. Yet, in most cases belief propagation lacks any performance and convergence guarantees. Many realistic problems are presented by graphical models with loops, however, in which case belief propagation is neither guaranteed to provide accurate estimates nor that it converges at all.

This thesis investigates how the model parameters influence the performance of belief propagation. We are particularly interested in their influence on (i) the number of fixed points, (ii) the convergence properties, and (iii) the approximation quality. For this purpose, we take a different perspective on belief propagation and realize that the fixed points define a set of polynomial equations – albeit a large one. Solving polynomial equations of this size is problematic; nonetheless, we present the numerical polynomial homotopy continuation method that is capable of solving the fixed point equations, the solutions of which are the fixed points of belief propagation.

The solutions to the fixed point equations gives us knowledge of the whole solution space and serves as a stepping stone for analyzing belief propagation’s properties. In particular, we observe a large variety of marginal accuracy across all fixed points. This, to some degree, explains the large discrepancy in the performance of belief propagation. Another important aspect of belief propagation’s fixed points is their stability, that is if belief propagation can – at least in principle – converge to a given fixed point. Existing stability analyses were limited to models without local potentials for the lack of knowing the solution space. The capability to solve the fixed point equations thus allows us to extend the stability analysis to a wide range of models. In doing so, we obtain novel insights into how the model parameters and the model size affect the stability. In particular, we find that strong pairwise potentials degrade the performance, whereas strong local potentials enhance the performance.

Moreover, our theoretical findings inspire a simple, yet powerful, modification of belief propagation. We present self-guided belief propagation that starts from a simple model (for which belief propagation obtains the exact solutions) and iteratively adapts it to the desired model. As the model is modified, self-guided belief propagation keeps track of the solution; this way, it improves upon standard belief propagation and it obtains the best possible solution for attractive models with unidirectional local potentials. For more general models, we empirically show that self-guided belief propagation maintains its favorable properties, converges more often, and is superior in terms of marginal accuracy.

Finally, we question whether the global minimum of the Bethe free energy provides the most accurate marginals. This is a common conjecture that inspired a range of methods that aim to minimize the Bethe free energy. In the past, the studied models were either too simplistic or too complex as to the true nature of this relationship. Therefore, we must first introduce a novel class of models – termed patch potential models – which are simple enough so that we can compute all fixed points. Yet, patch potential models are complex enough to possess a rich and non-trivial solution space. A study of this solution space proves this conjecture wrong and, additionally, explains the nature of the difference between accurate marginals and good approximations of the free energy.

Acknowledgments

There are a lot of people who influenced this thesis and made working on it much more enjoyable. Here I want to express my gratitude for all of you.

First of all, I must thank my supervisor Franz Pernkopf for being such a great advisor. You somehow always managed to find time for stimulating discussions, both about research and about life in general. Most importantly, however, you not only granted me the freedom to pursue my own ideas but you also believed in me when I struggled to do so.

Thank you, Bernhard and Gernot for supervising my Master's thesis. It is because of you I had such a good experience back then which made me seriously consider pursuing a PhD at this lab.

I want to thank all members of the SPSC I had the pleasure to meet during my time here – you are responsible for creating an awesome environment to work in. There are several people I want to thank in particular: Robert and Sebastian for the warm welcome I received when starting here, for the thought-provoking discussions, and for showing me how to survive in academia. Josef for being the best room-mate I can think of, for sharing the white-board with me, and for introducing me to space-exploration (as well as for discussing all kinds of research questions). Wolfgang for always having an open door and an open mind (and a good supply of sweets). Johannes for our endless discussions about research, teaching, and – most importantly – everything else. Thomas for making sure that our kitchen is always well-stocked with the essentials. Stefan, Erik, Jamilla, Michi, Martin T., Elmar, Johanna and Alex for regularly testing the capacity of our coffee-machine – you certainly enriched the time I spent standing in the queue. I am pretty sure I forgot to mention some of you – please forgive me and do not take this personally!

Looking back at how I got here, I must thank those who probably shaped me the most during my whole life: A big thank you to my parents and my sister. Thanks for your support, for always being honest, and for listening to me whenever I need someone to talk to.

Finally, I must thank my own small family. Thank you, Jakob and Hanna, for constantly reminding me that there are things in life far more important than my research. And, of course, I own my deepest gratitude to you, Kati! Thank you for your endless support; thank you for sometimes understanding me better than I do myself; and thank you for learning me how to appreciate and enjoy even the little success-stories in life.

Notational Conventions

General Notation

\mathbb{R}_+^*	positive real numbers, excluding zero
\mathbb{Z}_+^*	positive integer numbers, excluding zero
∇	gradient
∇^2	Hessian
$D(\tilde{P}_{\mathbf{X}} P_{\mathbf{X}})$	Kullback Leibler divergence between $\tilde{P}_{\mathbf{X}}$ and $P_{\mathbf{X}}$
$\text{Re}\{\cdot\}, \text{Im}\{\cdot\}$	real and imaginary part
i	imaginary unit
$ \cdot $	magnitude or absolute value
\oplus	exclusive disjunction
$\text{sgn}(\cdot)$	sign function
$n \pmod{m}$	n modulo m
$\mathbf{1}$	indicator function
$\log(\cdot)$	natural logarithm
$\ \cdot\ _p$	l_p -norm
$\mathcal{N}(\mu, \sigma^2)$	Gaussian distribution with mean μ and variance σ^2
$\mathcal{U}(a, b)$	uniform distribution on (a, b)

Probability

X, Y, \dots	random variables
$\mathcal{X}, \mathcal{Y}, \dots$	range of a discrete random variable
x, y, \dots	value of random variable
\mathbf{X}	set of random variables
\mathbf{x}	values for a set of random variables (configuration)
\mathcal{X}^N	product space, range of \mathbf{X}
$P_X(x)$	probability distribution of X
$\tilde{P}_X(x)$	approximation of the probability distribution of X
\tilde{P}_B	pseudomarginals (set of all singleton and pairwise marginals)
Z_i, Z_{ij}	normalization terms for approximated marginals
$P_{X Y}(x y)$	conditional probability distribution of X given Y
$P_{\mathbf{X}}(\mathbf{x})$	Joint probability distribution of \mathbf{X}
$\mathbb{E}(x)$	expected value of X
m_i	mean of X_i
$\langle m \rangle$	expected mean of \mathbf{X}
χ_{ij}	correlation between X_i and X_j

Graphs

\mathcal{G}	undirected graph
X_1, X_2, \dots	nodes
\mathbf{X}	set of nodes
N	number of nodes
(i, j)	edge between X_i and X_j
m	number of the edge (i, j) under some ordering
\mathbf{E}	set of edges
$\partial(i)$	neighbors of X_i
d_i	degree of X_i
$\hat{d}_{\mathcal{G}}$	average degree of \mathcal{G}
\mathcal{G}'	subgraph of \mathcal{G} induced by \mathbf{X}'
\mathbf{A}	adjacency matrix

Graphical Models

\mathcal{U}	undirected graphical model
C_i	clique
$C_{(i,j,\dots)}$	clique consisting of X_i, X_j, \dots
\mathbf{C}	set of cliques
$\Phi_{C_i}(\mathbf{x}_{C_i})$	clique-potential
Ψ	set of potentials
$\Phi_{X_i}(x_i), \Phi(x_i)$	local potential
$\Phi_{X_i, X_j}(x_i, x_j), \Phi(x_i, x_j)$	pairwise potential
\mathcal{Z}	partition function
$E(\mathbf{x})$	energy of the configuration \mathbf{x}
J_{ij}	coupling
θ_i	local field
σ_i	spin
$\boldsymbol{\sigma}$	configuration of spins
β	inverse temperature
H	external magnetic field
$J_{\mathcal{P}}$	product of couplings along path \mathcal{P}

Dynamical Systems and Equation-Systems

\mathbf{x}	state vector
$\mathbf{x}(n), \mathbf{x}^n$	state vector \mathbf{x} at time index n
\mathbf{x}°	fixed point of \mathbf{x}
$U(\mathbf{x})$	ϵ -neighborhood of \mathbf{x}
$\mathcal{F}(\cdot)$	discrete-time map
\mathbf{B}	system matrix
$\mathcal{F}'(\mathbf{x})$	Jacobian matrix
λ_i	eigenvalue
λ_{max}	eigenvalue with the largest magnitude
$\Lambda(\mathbf{A})$	set of all eigenvalues for the matrix \mathbf{A}
$\rho(\mathbf{A})$	spectral radius of the matrix \mathbf{A}
$f_i(\mathbf{x})$	function or polynomial equation
a_k	polynomial coefficient
$\mathbb{K}[x_1, \dots, x_n]$	polynomial ring
$\mathbf{F}(\mathbf{x}), \mathbf{F}(x_1, \dots, x_n)$	system of polynomial equations
$\mathbf{V}(\mathbf{F})$	variety (set of solutions) of $\mathbf{F}(\mathbf{x})$
$\mathbf{V}_{\mathbb{R}}(\mathbf{F}), \mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$	variety of $\mathbf{F}(\mathbf{x})$ for the real and the positive real numbers
S_i	convex hull of the exponent vectors of f_i
$V(S_i)$	volume of the polytope S_i
$M(S_i, S_j)$	mixed volume of S_i and S_j
ω_i	lifting for the polynomial f_i
$\omega_i(a)$	lifting function for the exponent-vector a
$\langle f_1, \dots, f_s \rangle$	ideal of \mathbf{F}
d_t	total degree
BKK	Bernshtein-Kushnirenko-Khovanskii bound
$\mathbf{H}(\mathbf{x}, t)$	homotopy
$\mathbf{Q}(\mathbf{x})$	start system
t	goes from 0 to 1 to deform the homotopy
γ	random complex number

Belief Propagation, Variational Approximations

$\mu_{ij}^n(x_j)$	message from X_i to X_j
$\boldsymbol{\mu}^n$	set of messages
n	iteration index
α_{ij}^n	message normalization for $\mu_{ij}^n(x_j)$
$\mu_{ij}^\circ(x_j)$	fixed point message
$\boldsymbol{\mu}^\circ$	fixed point of belief propagation
$E(\tilde{P}_{\mathbf{X}})$	average energy
$S(\tilde{P}_{\mathbf{X}})$	entropy
$\mathcal{F}_{\mathcal{H}}, \mathcal{F}_{\mathcal{G}}, \mathcal{F}_{\mathcal{B}}$	Helmholtz-, Gibbs-, and Bethe- free energy
$E_{\mathcal{B}}(\tilde{P}_B)$	average (Bethe) energy
$S_{\mathcal{B}}(\tilde{P}_B)$	Bethe entropy
$\mathbb{M}(\mathcal{G}), \mathbb{M}$	marginal polytope
$\mathbb{L}(\mathcal{G}), \mathbb{L}$	local polytope
\mathcal{Z}_B	Bethe partition function
$\mathcal{F}_{\mathcal{B}}^*, \mathcal{F}_{\mathcal{B}}^m, \mathcal{F}_{\mathcal{B}}^\circ$	global minimum, local minimum, and stationary point of $\mathcal{F}_{\mathcal{B}}$
$\tilde{P}_B^*, \tilde{P}_B^m, \tilde{P}_B^\circ$	pseudomarginals at the global minimum, a local minimum, and a stationary point of $\mathcal{F}_{\mathcal{B}}$
S	set of stable belief propagation fixed points
M	set of fixed points corresponding to local minima of $\mathcal{F}_{\mathcal{B}}$
T	set of all fixed points
$E_Z(m)$	error of the partition function approximation
$E_P(m)$	error of the approximated marginals
$\mathcal{BP}(\cdot)$	mapping induced by belief propagation
\mathcal{D}_P	evaluation function for the pseudomarginals
\mathcal{D}_Z	evaluation function for the Bethe partition function
r_{ij}^n	message-residual
\mathbf{r}^n	set of residuals
ϵ	damping factor of belief propagation with damping
$\mathbf{F}(\boldsymbol{\mu}, \boldsymbol{\alpha})$	fixed point equations of belief propagation
ν_{ij}^n	reparameterized messages
$\boldsymbol{\nu}$	set of reparameterized messages
h_{ij}	cavity field
\tilde{P}_B^T	weighted combination of all fixed points
\tilde{P}_B^M	weighted combination of all fixed points belonging to local minima
\tilde{P}_B^{MAX}	fixed point maximizing the partition function
$J_A, J_C, J_C(\mathcal{G}, \theta)$	critical values of coupling strength at the onset of phase-transitions
κ_θ	field-dependent scaling term for Jacobian matrix

Error Correcting Codes

ϵ	error-probability
Y_i	variable node
f_a	factor node
$r_{A_i}^n(y_i)$	message from factor to variable
$q_{iA}^n(y_i)$	message from variable to factor
α_{iA}^n	normalization term

Self-Guided Belief Propagation

ζ_k	scaling term
K	number of models considered by self-guided belief propagation (length of $\{\zeta_k\}$)
\mathcal{U}_K	undirected graphical model for ζ_k
Ψ_k	set of potentials for ζ_k
$\Phi_k(x_i, x_j), \Phi_k(x_i)$	pairwise and local potentials for ζ_k
μ_k°	fixed point of BP for the model \mathcal{U}_k
$c(\zeta)$	solution path
N_{BP}	maximum number of belief propagation iterations

Patch Potential Models

\mathcal{G}_i	patch of the graph \mathcal{G}
P_μ^m	fraction of runs converging to the m^{th} fixed point
$\tilde{P}_B^r, \tilde{P}_B^q$	fixed points with all marginals biased to one state
\tilde{P}_B^p	state-preserving fixed point
$\tilde{P}_B^u, \tilde{P}_B^v, \dots$	all other fixed points
$\tilde{\theta}_i$	effective field for X_i
$Q_i(k, l)$	mismatch between \tilde{P}_B^k and \tilde{P}_B^l
\mathbf{E}_P	set of all boundary edges
\mathbf{E}_C	set of edges between variables that favor different states
N_f	number of flipped variables
N_c	number of state-preserving variables
ΔS_B	difference of the Bethe entropy between two fixed points

1

Introduction

*”Uncertainty is an uncomfortable position.
But certainty is an absurd one.”*
– Voltaire

1.1 Motivation

Every realistic domain, in accordance with the opening quote, contains some degree of uncertainty: not only is every human imperfect – with each of our decisions being subject to uncertainty – but also has every artificial system only access to partial information. When reasoning under such uncertainties, one must take all available information into account and infer the quantity of interest. This task of probabilistic reasoning is one of the central problems of statistical inference¹ and, additionally, lies at the center of every decision-making process.

Most inference-problems of practical relevance belong to the realm of multivariate statistics and involve many interacting variables. As the number of variables increases, the corresponding distributions become increasingly hard to grasp and thus require compact representations; particularly if one wishes to retain some interpretability. Fortunately, such a representation exists in the form of probabilistic graphical models; these models (as suggested by their name) rely on the proven capabilities of graphs² and represent complex interactions in an intuitive and expressive way.

We will study two of the most fundamental problems of inference. These are: computing the *marginal distribution* and evaluating the *partition function*. Both problems suffer from the increased complexity when working with many variables. Assistance comes in the form of graphical models once again; not only do they represent the problems efficiently, but they also warrant efficient inference algorithms that exploit the graph structure and thus facilitate the task of inference. Note, however, that graphical models are not a panacea and while inference becomes tractable for some models (e.g., trees) it remains NP-hard for graphs with loops.

Besides being of theoretical interest, such loopy graphs arise in a multitude of practical applications; ranging from statistical signal-, speech-, and image-processing, to statistical physics, medical diagnosis systems, and error-correcting codes. For all these important problems, exact inference methods, however, are destined to fail, which substantiates the need for efficient approximation methods.

One particularly prominent approximation method focuses on local interactions and infers the global model-behavior thereof. Consider, for example, a large group of friends that want to celebrate a party and need to find a suitable date. If all friends come together and discuss their

¹ One may argue that statistical inference provides the common ground for a myriad of scientific fields ranging from mathematics over empirical sciences to philosophy [EH16]. We comply with this perspective and provide a broad context to the rather specific topic of this thesis.

² Graphs naturally emerge in a wide range of problems that include applications in social sciences, biology, physics, and communications. Graph theory constitutes a concept with many persuasive properties that warrant their wide-spread usage. First, graphs lend themselves for visual representations that reveal the underlying problem-structure, often hidden initially, to the human observer. Second, manipulations and computations on graphs are well established and provide an extensive tool-box to tackle a problem once formulated as graph.

preferences in one large negotiation, they would find an optimal date; but this seems overly complicated and impracticable. Instead, one could hope to find an acceptable date that works for most by negotiating in a distributed manner: assume that everybody discusses his preferences only with his closest friends (and that these sub-groups sufficiently overlap), then a reasonable agreement on the date will be found that will suit at least the majority of the group. Note how this focus on local interactions comes relatively intuitive when dealing with large, complex systems. The graphical model's structure is exploited precisely in such a way by belief propagation that performs local interactions – in the form of exchanging messages between neighboring nodes – to approximate the marginal distribution and the partition function.

Let us briefly summarize the main strength of belief propagation (BP); it is the ability to efficiently perform approximate inference on models where (because of loops) exact inference becomes infeasible. BP often works remarkably well in this context, although theoretical results fail to explain BP's empirical success. Yet, portraying BP as an outright success-story would not correspond to the truth either – and indeed, while BP often performs approximate inference in a very efficient manner it completely fails to do so other times. For this lack of reliability, BP is sometimes confronted with skepticism. Instead of only hoping for a reasonable performance of BP, we would, ideally, like to have some performance guidelines established. Stating whether a given model is well-suited for the application of BP, such guidelines must focus on the following two aspects.

First, one needs to understand the underlying reasons for the failure of BP. Although we still lack a rigorous understanding of why BP fails, it is empirically well-established that BP fails because: (i) multiple solutions exist with varying accuracy; (ii) one or all fixed points are unstable so that BP does not converge.

Second, one needs to derive performance guarantees that take the model specifications into account. If we want to better understand BP we will not only need to understand how the model specifications influence (i) the convergence properties and (ii) the approximation quality but also how both properties relate to each other. Note that both properties are directly related to each other for simple models (e.g., graphs with a single loop or small grid graphs), i.e., the better the approximation quality the faster BP converges [Wei00, Ihl07]. Such a relation, however, does not generalize to more complex models [WJ14a].

To summarize, it will be important to understand how specific models affect the performance of BP and if BP can be expected to perform well. This will also increase the reliability of BP.

To a large degree, we owe our current understanding of BP to concepts from statistical physics. As it turns out, there is a fundamental connection between many concepts in computer science and in physics (cf. [MM09, WT03, TJ02]). Most notably, the fixed points of BP are in a one-to-one correspondence with the stationary points of the Bethe free energy.

But, although the Bethe free energy can provide many insights, it remains an intricate function that is hard to analyze. Also, most theoretical results on the Bethe free energy only hold for restricted model classes, where typically all variables sit on a regular grid and the model is specified by relatively few parameters. How those insights carry over to more general models is an open question.

Moreover, the Bethe free energy fails to reveal whether a given fixed point is stable and under which conditions BP converges to it. The Bethe free energy also fails to explain why and how certain modifications of BP (e.g., scheduling or damping) often help to achieve convergence.

The quest for a better understanding of BP is thus ongoing, with the hope that new insights will recognize certain model classes for which it is safe to utilize BP, i.e., for which BP converges fast while maintaining the desired accuracy.

The overarching aim of this thesis is to extend the current understanding of BP, the reason for this being twofold: first, to theoretically understand for which problems and applications BP

can be expected to perform well; and second, to utilize those theoretical insights and modify BP to enhance its capabilities.

We specifically address the question of how knowledge of the solution space can advance our current understanding of BP. As the complete set of solutions is generally not available, we must first develop a way to obtain all fixed points for a given model. We then start with the analysis of relatively small and simple models and, by successively building upon our obtained insights, extend our analysis to increasingly more complex models. Finally, we take a comprehensive view at the solution space and study the relation between the convergence properties, approximation quality, and the number of fixed points. This approach will provide several insights into the behavior of BP, extend the current theoretical understanding and open the door for practical considerations that enhance the performance of BP. In particular, our findings suggest a modification of BP that enforces convergence toward accurate fixed points, thus improving the approximation quality.

1.2 Five Relevant PhD Theses

As discussed, BP touches a diverse set of scientific fields. It is the common ground between all those fields that provides a solid foundation for the analysis of BP. The work presented in this thesis therefore builds upon a vast body of literature. The following five PhD theses cover a wide range of the recent developments and are highly relevant for the current thesis. Therefore, this section contains a brief overview of how they shaped the current thesis in particular.

The connection between exponential representations of distributions, information geometry, and approximate inference methods on probabilistic graphical models is revealed in the thesis of **Martin Wainwright**, submitted at the Massachusetts Institute of Technology in 2002 [Wai02]. His thesis introduces the concept of reparameterization. This concept casts BP as one specific instance in a more general class of message passing algorithms (termed tree-based reparameterization (TRP)). Moreover, TRP suggests solving a particular sequence of simpler sub-problems over spanning trees in the graph [WJW03a] in order to enhance the convergence properties. Additionally, important insights with respect to the marginal accuracy are obtained in the form of an exact expression for the marginal error, that – although being infeasible to evaluate in general – suggests computable bounds on the marginal error. The proposed bounds rely on the approximation of the log-partition function which advocates a close connection between both quantities; this connection nicely connects to the present work where we inspect this relationship in great detail in Chapter 7. Another powerful concept is the consideration of BP as an optimization problem (over the local polytope). In the conclusion, the author proposes tracing the evolution of the pseudomarginals while relaxing the marginal- to the local polytope with the prospect of practical and theoretical consequences; a similar evolution of the pseudomarginals lies at the core of Chapter 6.

The overarching aim of the thesis of **Joris Marten Mooij**, submitted at the Radboud Universiteit Nijmegen in 2008 [Moo08], is to understand and improve belief propagation, which closely resembles the aim of the current thesis. In a nutshell, his thesis investigates the relationship between accuracy, uniqueness, and convergence properties of belief propagation’s fixed points. One particular insightful contribution was the consideration of belief propagation as a dynamical system. This led to conditions for stability and uniqueness of a fixed point, as well as insights into the relation between those properties. The analysis, however, was restricted to vanishing local potentials. We adhere to the spirit of considering belief propagation as a dynamical system and extend the analysis to models with arbitrary parameters in Chapter 5, providing novel theoretical insights.

The application of belief propagation to inference problems arising in the context of sensor

networks is the main motivation in the thesis of **Alexander Ihler** [Ihl05], submitted at the Massachusetts Institute of Technology in 2005. His thesis focused on studying the fundamental limitations of belief propagation. Studying the vulnerability of the approximation quality with respect to errors in the messages was one particularly important aspect with the prospect of validating whether approximating the messages is a viable option. Notably, this leads to a couple of interesting theoretical results on the accuracy of belief propagation in general, and the introduction of error-bounds on the marginal accuracy specifically. While these bounds are only valid for cycle-free graphs, they predict the performance of belief propagation reasonably well in the presence of loops as well. This is one of the few results that analyze the error in the marginals, while much of the literature focuses on the approximation error of the partition function. We agree with the author on the importance of assessing the marginal accuracy and will particularly focus on how the model parameters influence the marginal accuracy (cf. Chapter 5 and 7).

The thesis [Well14] of **Adrian Weller**, submitted at the Columbia University in 2014, focuses on the variational interpretation of belief propagation and the Bethe approximation in particular. His work provides interesting insights into the differences between log-partition function estimates, singleton marginals, and pairwise marginals for both the (non-convex) Bethe approximation and convex variational approaches. In particular, it becomes evident that different approximation methods do not affect the marginals and the log-partition function in the same way; the accuracy of the log-partition function may for example remain the same whereas the accuracy of the marginals increases. This observations raises the question of how both quantities are related, an important question that will be the main focus of Chapter 7. Besides the theoretical relevance of the obtained insights, his thesis further proposes an approximation of the Bethe function, which extends preceding work [Shi12] by approximating the global minimum. This approximation method converges in polynomial runtime for attractive models and serves as an important comparison for our proposed method in Chapter 6.

A direct relationship between belief propagation and graph geometry is established in the thesis of **Yusuke Watanabe** [Wat10], submitted at The Graduate University for Advanced Studies, SOKENDAI in 2010; this relationship provides novel insights into the behavior of belief propagation. In particular, the *graph-zeta function* relates the convergence properties of belief propagation to the shape of the Bethe free energy.³ Moreover, the results impose some general properties on the solution space. These properties are of relevance for belief propagation as well and, for example, demonstrate that the overall number of fixed points is always odd. Computing the number of fixed points is also of central interest throughout the current thesis, and Chapter 5 and 7 in particular.

1.3 Contribution and Outline

A large part of the contributions to this thesis has previously been published; a list of the corresponding publications is presented below. Several parts, however, have been significantly reworked and restructured in order to align nicely with the structure of the thesis.

- [KRTP15] Christian Knoll, Michael Rath, Sebastian Tschitschek, and Franz Pernkopf. Message scheduling methods for belief propagation. In *Proceedings of ECML PKDD*, pages 295–310. Springer, 2015.
- [KPMC16] Christian Knoll, Franz Pernkopf, Dhagash Mehta, and Tianran Chen. Fixed point solutions of belief propagation. In *NIPS-Workshop: Advances in Approximate Bayesian Inference*, 2016.

³ In addition to demonstrating the wide range of fields that play an important role for our current understanding of belief propagation, the graph-zeta function also serves as a foundation for recent developments in spectral clustering [SKZ14].

- [KP17] Christian Knoll and Franz Pernkopf. On loopy belief propagation – local stability analysis for non-vanishing fields. In *Proceedings of UAI*, 2017.
- [KMCP18] Christian Knoll, Dhagash Mehta, Tianran Chen, and Franz Pernkopf. Fixed points of belief propagation – an analysis via polynomial homotopy continuation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [KKP18] Christian Knoll, Florian Kulmer, and Franz Pernkopf. Self-guided belief propagation – a homotopy continuation method. *arXiv preprint arXiv:1812.01339*, 2018.
- [KP19] Christian Knoll and Franz Pernkopf. Belief propagation: Accurate marginals or accurate partition function – where is the difference? In *Proceedings of UAI*, 2019.

It is the very nature of most PhD theses to pinpoint and raise multiple questions before providing some – hopefully insightful – answers to them. As is often the case, this thesis tackles very specific problems in an already highly specialized field. One aspect, however, that widens the scope and that made working on this thesis particularly interesting is that belief propagation is applied in various scientific fields ranging from information theory and signal processing to statistical physics and artificial intelligence. Our study of belief propagation brings insights from all these fields together and benefits from a particularly large toolbox, fueled by such diverse inputs. While it is interesting to delve into each particular field and emerge oneself in all the associated details, it is nearly impossible to hide one’s past. Therefore, and for the purpose of a consistent notation, we settle for the language used in the statistics and machine learning community, point at relations to associated scientific fields if appropriate, and draw from them if beneficial.

Writing a coherent thesis, even after completing the research tasks, remains an extensive task that should ideally serve the interested reader. Therefore, this thesis – rather than being a conglomerate of the presented results – offers a thorough introduction and sticks to one coherent story. In doing so, we identify an inherent structure that lends itself to a segmentation into two major parts.

The first part (Chapter 2 - 4) provides all relevant background and serves as the preparation for the subsequent chapters.

The whole thesis resides in the context of probabilistic graphical models, which are introduced in **Chapter 2**. We briefly discuss the relevant background from probability- and graph-theory and introduce pairwise graphical models that are the main focus of this work. Moreover, we define the problem of inference with a particular focus on efficient exact methods.

Belief propagation (BP) is introduced as a method of approximate inference in **Chapter 3**. After discussing some of the most serious issues of BP, we describe some of the underlying reasons for failure of BP and introduce alternative characterizations of BP. In particular, we introduce BP as a variational method and connect the properties of BP with the energy landscape of the Bethe free energy. Additionally, we cast BP as a dynamical system. It turns out that this characterizations provides a general framework, encompassing a wide range of approximate inference methods. Moreover, the consideration as a dynamical system suggests various ways to enhance the properties of BP.

Chapter 4 provides the reader with the most important background of dynamical system theory and algebraic geometry. While this chapter does not contain novel insights, it provides us with the relevant tools for analyzing BP in detail. In particular, the most prominent methods for solving system of equations are introduced and discussed.

The second part (Chapter 5-7) builds upon the knowledge developed so far, presents the major results and extends the current understanding of BP’s behavior.

The consideration of BP as a dynamical system provides the foundation of the results developed in **Chapter 5**. We apply tools from dynamical systems theory, analyze the solution space of BP, and gain new insights into the behavior of BP; the focus on the whole solution space instead of just a single fixed point reveals how the number of fixed points, the approximation quality of the individual fixed points, and the convergence properties are related to each other. One major issue is that, despite being conceptually straightforward, it is problematic to find the set of *all* fixed points. We present how the set of all BP fixed points can be computed by using the numerical polynomial homotopy continuation (NPHC) method. This allows us to assess and compare the accuracy of the individual BP fixed points and weighted combinations thereof. Moreover, the knowledge of all fixed points allows us to extend the local stability analysis – previously restricted to models with vanishing local potentials – to more general models. This generalization also explains the role of the local potentials and reveals how strong local potentials enhance the convergence properties.

The focus of **Chapter 6** is to exploit the theoretical finding that local and pairwise potentials play an opposing role regarding the performance of BP. In this chapter we propose one way to account for this observation and to enhance the performance of BP. By modifying BP according to a homotopy continuation method we account for this observation and incorporate the pairwise potentials only gradually. This procedure is deterministic, converges to a uniquely defined fixed point, and thus resolves the dependence on a well-chosen initialization. Experiments on a wide range of models reveal that the proposed method increases the performance without increasing the computational burden; in particular we exemplify that the results are at least as accurate as BP, if BP converges, and that accurate results are often obtained, even if BP fails to converge. This empirical analysis is further supplemented with a theoretical analysis that proves optimality – i.e., the proposed method converges to the global minimum of the Bethe free energy that constitutes the fixed point with the most accurate marginals – for restricted models that have all local potentials favoring the same state.

All models considered so far adhere to the common conjecture that accuracy with respect to the marginals and with respect to the partition function are interchangeable. This conjecture suggests that the most accurate marginals of BP are to be found at the global minimum of the Bethe free energy. In **Chapter 7** we aim to validate this assumption. We therefore introduce *patch potential models* that simplify the analysis significantly. The parameter space of patch potential models is split into multiple regions with fundamentally different properties. Elaborating on one specific region we first exemplify why there is no strict relationship between the accuracy of the marginals and the partition function and then provide sufficient and necessary conditions for a fixed point to be optimal with respect to approximating both.

Finally, **Chapter 8** concludes this thesis, summarizes the results obtained, and discusses the extent to which our understanding of BP has changed. Furthermore, the most pressing questions left unanswered are indicated, hence providing a pointer to potential future research directions.

2

Background

*"A mind is like a parachute.
It doesn't work if it is not open."
– Frank Zappa*

This chapter introduces the necessary background and lays the foundation for the subsequent chapters. The focus of this thesis lies at problems that arise in the context of probabilistic graphical models. Accordingly, we begin with a brief introduction to probability theory, in Section 2.1, and to graph theory, in Section 2.2, before we finally unite both concepts in probabilistic graphical models, in Section 2.3. Section 2.4 discusses one specific class of probabilistic graphical models, *undirected binary pairwise models*, that are studied extensively throughout the thesis. We finally devote Section 2.5 to the tasks summarized under the notion of *inference*.

2.1 Probability Theory

We begin with a brief review of probability theory and introduce the most relevant notions. Note that we will restrict our focus to discrete random variables throughout this thesis. Let us therefore consider a discrete random variable X that maps from the sample space Ω to a discrete finite set \mathcal{X} , i.e., $X : \Omega \rightarrow \mathcal{X}$.

We denote the probability mass function that assigns a probability to the generic value $x \in \mathcal{X}$ by the shorthand notation $P_X(x) = P_X(X = x)$. Likewise, let X and Y be two discrete random variables; then, with slight abuse of notation we define the joint distribution according to

$$P_{X,Y}(x, y) = P_{X,Y}((X = x) \cap (Y = y)). \quad (2.1)$$

Now let us consider a set of N discrete random variables $\mathbf{X} = \{X_1, X_2, \dots, X_N\}$ with the joint distribution $P_{\mathbf{X}}(\mathbf{x})$ where the range of \mathbf{X} is the product space $\mathcal{X}^N = \mathcal{X}_1 \times \dots \times \mathcal{X}_N$. We further denote the configuration for a given set of random variables by $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$. The marginal probability distribution for a subset $\mathbf{Y} \subset \mathbf{X}$ is obtained by summing out all variables X_i that are not in \mathbf{Y} according to

$$P_{\mathbf{Y}}(\mathbf{y}) = \sum_{X_i \in \{\mathbf{X} \setminus \mathbf{Y}\}} \sum_{x_i \in \mathcal{X}_i} P_{\mathbf{X}}(\mathbf{x}). \quad (2.2)$$

(Conditional) statistical independence between random variables is a particularly important property when dealing with probabilistic graphical models. We say that \mathbf{X} is conditionally independent of \mathbf{Y} given \mathbf{Z} whenever $P_{\mathbf{X},\mathbf{Y}|\mathbf{Z}}(\mathbf{x}, \mathbf{y}|\mathbf{z}) = P_{\mathbf{X}|\mathbf{Z}}(\mathbf{x}|\mathbf{z})P_{\mathbf{Y}|\mathbf{Z}}(\mathbf{y}|\mathbf{z})$ holds for all configurations $\mathbf{x} \in \mathcal{X}^N$, $\mathbf{y} \in \mathcal{Y}^N$, and $\mathbf{z} \in \mathcal{Z}^N$.

Distributions are often characterized by their expectations. For a discrete random variable X we define its expectation under the distribution $P_X(x)$ according to

$$\mathbb{E}(x) = \sum_{x \in \mathcal{X}} x \cdot P_X(x). \quad (2.3)$$

2.2 Graph Theory

One of the most persuasive properties of probabilistic graphical models is the representation of joint probability distributions by using graphs. This section gives a self-contained introduction to the basics of graph theory. We refer the interested reader to one of the many available books on graph theory for a more in-depth treatment (e.g., [KV05, Die16]).

Let us consider a *graph* $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ with a set of *nodes* (or vertices) $\mathbf{X} = \{X_1, \dots, X_N\}$ and a set of *edges* \mathbf{E} . A graph is either directed or undirected and consists of directed or undirected edges respectively. We denote an undirected edge by $(i, j) \in \mathbf{E}$ (or equivalently by (j, i)) if it joins two nodes $X_i \in \mathbf{X}$ and $X_j \in \mathbf{X}$. Note that, for the remainder of this thesis, we will only consider undirected graphs and further restrict our focus to *simple graphs* – that come without parallel edges (that would join the same pair of nodes twice) and without self-loops (i.e., $i \neq j$).

Let us consider an edge $(i, j) \in \mathbf{E}$, then X_j is a *neighbor* of X_i and vice versa; the set of neighbors for any variable $X_i \in \mathbf{X}$ is defined by

$$\partial(i) = \{X_j \in \mathbf{X} : (i, j) \in \mathbf{E}\}. \quad (2.4)$$

The *degree* of a node X_i is the number of incident edges and is consequently defined by the cardinality of the neighbor-set, i.e., by

$$d_i = |\partial(i)|. \quad (2.5)$$

We further denote the *average degree* of \mathcal{G} by $\hat{d}_{\mathcal{G}} = \frac{2|\mathbf{E}|}{N}$. A *path* between two nodes $X_i, X_k \in \mathbf{X}$ is a graph $\mathcal{P} = (\{X_i, \dots, X_k\}, \{(i, i+1), \dots, (k-1, k)\})$, where X_i and X_k are the end-nodes of \mathcal{P} . A path goes from X_i to X_k if there is an edge-progression that connects those two nodes.

Definition 2.2.1 (Loop). A loop is a path from a node X_i back to the same node along a sequence of edges, i.e., a path $\mathcal{P} = (\{X_i, X_{i+1}, \dots, X_k\}, \{(i, i+1), \dots, (k-1, k)\})$ exists for $k = i$.

Depending on the notational conventions, loops are sometimes also referred to as cycles or circuits. Note that every loop has to contain at least three edges, since we do not allow for self-loops.

A graph is said to be *connected* if, for every pair of nodes X_i and X_j , a path exists that connects X_i to X_j . We only consider connected graphs in this thesis.

In the context of this thesis, there are certain graph types that are particularly relevant; most notably these include the following:

Definition 2.2.2 (Tree). A connected graph that does not contain any loops is a tree.

Definition 2.2.3 (Complete Graph). A complete graph $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ has every pair of nodes X_i and X_j connected by an edge (i, j) ; i.e., \mathcal{G} is fully connected with $\mathbf{E} = \{(i, j) : X_i, X_j \in \mathbf{X}, i \neq j\}$. It follows that every node has an equal degree $d_i = N - 1$. See Figure 2.1 for an illustration of complete graphs of size one to four.

Definition 2.2.4 (Grid Graph). An $m \times n$ grid graph, or lattice graph, is a graph with $N = mn$ nodes that has all edges aligned along the square lattice. See Figure 2.2 for an illustration of a two-dimensional grid graph.

It will often be necessary to consider only a part of the graph; we call this a subgraph.

Definition 2.2.5 (Subgraph). Let $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ be a graph and let $\mathbf{X}' \subseteq \mathbf{X}$. Then, the subgraph (or graph-component) $\mathcal{G}' = (\mathbf{X}', \mathbf{E}')$ is induced by \mathbf{X}' , where $\mathbf{E}' = \{(i, j) \in \mathbf{E} : X_i, X_j \in \mathbf{X}'\}$.




Figure 2.1: Complete graphs of size $N = 1$ to $N = 4$.




Figure 2.2: Two-dimensional grid graph of size $N = 3 \times 3$

Complete subgraphs are of particular relevance and are often referred to as cliques.

Definition 2.2.6 (Clique). Let $\mathcal{G}' = (\mathbf{X}', \mathbf{E}')$ be a subgraph. If \mathcal{G}' is complete we call it a clique and refer to it as C_i . If any adjacent node $X_j \in \mathbf{X} \setminus \mathbf{X}' : (i, j) \in \mathbf{E}, X_i \in \mathbf{X}'$ exists that, by adding, would render C_i not complete anymore, C_i is a maximal clique.

See Figure 2.1 that depicts the maximal cliques of size one to four. Note that any graph $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ is complete if and only if \mathcal{G} is a maximal clique.

Definition 2.2.7 (Regular Graph). Let $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ be a graph where all nodes $X_i \in \mathbf{X}$ have equal degree $d_i = d$. Then, we refer to \mathcal{G} as a d -regular graph.

It follows that the average degree for a regular graph equals the degree of all individual nodes, i.e., $\hat{d}_{\mathcal{G}} = d_i$.

Definition 2.2.8 (Bipartite Graph). A bipartite Graph $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ is a graph that decomposes into two disjoint subgraphs \mathbf{Y} and \mathbf{Z} where $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z}$ and $\mathbf{Y} \cap \mathbf{Z} = \emptyset$ so that every edge connects those two subgraphs; i.e., for all $Y_i \in \mathbf{Y}$ we have $\partial(Y_i) \in \mathbf{Z}$ and vice versa.

Definition 2.2.9 (Adjacency Matrix). The adjacency matrix \mathbf{A} represents the connections of a finite graph. \mathbf{A} is a 0 – 1 matrix with rows and columns indexed by the set of nodes so that $a_{ij} = 1$ if and only if $(i, j) \in \mathbf{E}$. Note that \mathbf{A} is symmetric for undirected graphs.

2.3 Probabilistic Graphical Models

Probabilistic graphical models provide a compact representation of joint distributions and are particularly well suited for representing distributions with many random variables. The straightforward specification of a distribution in N random variables requires one to define and store the probabilities of all $|\mathcal{X}|^N$ configurations; considering the fact that typical problems often have

hundreds of random variables renders such an approach impracticable. The representation of the joint distribution by the means of a graph on the other hand is intuitive and comes with the advantage of being interpretable by humans. Moreover, and of even greater relevance, the graph exploits the statistical dependencies of the distribution and makes them explicit. This is required if confronted with (even moderately) high-dimensional joint distributions.

Probabilistic graphical models come in different forms that represent the statistical dependencies in slightly different ways. These include the most prominent types such as factor graphs, Bayesian networks, and Markov random fields. We will, however, restrict our focus to Markov random fields (which are also termed undirected graphical models).

2.3.1 Undirected Graphical Models (Markov Random Fields)

A probabilistic graphical model $\mathcal{U} = (\mathcal{G}, \Psi)$ consists of an undirected graph $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ and a set of K potentials (sometimes called clique potentials, or compatibility functions) Ψ that defines the joint distribution $P_{\mathbf{X}}(\mathbf{x})$.

There is a certain elegance to graphical models; one that relates the structural properties of the graph to the properties of the associated joint distribution. Specifically, a one-to-one correspondence between the set of nodes $\mathbf{X} = \{X_1, \dots, X_N\}$ and the set of random variables⁴ holds. Additionally, each edge (i, j) represents the existence of a statistical dependency between X_i and X_j .

The correspondence between edges and statistical dependencies manifests itself into the global Markov property. Let $\mathbf{U} \subset \mathbf{X}$ be a subset of nodes; we say that \mathbf{U} separates two disjoint (sets of) nodes \mathbf{V} and \mathbf{W} (i.e., $\mathbf{V} \cap \mathbf{W} = \emptyset$) if every path connecting the sets \mathbf{V} and \mathbf{W} contains at least one variable $U_i \in \mathbf{U}$. Removing the subgraph induced by \mathbf{U} from \mathcal{G} thus eliminates all paths between \mathbf{V} and \mathbf{W} . Let us assume that \mathbf{U} separates \mathbf{V} and \mathbf{W} as discussed for \mathcal{G} ; then a joint distribution $P_{\mathbf{X}}(\mathbf{x})$ satisfies the global Markov property with respect to \mathcal{G} if

$$P_{\mathbf{V}, \mathbf{W} | \mathbf{U}}(\mathbf{v}, \mathbf{w} | \mathbf{u}) = P_{\mathbf{V} | \mathbf{U}}(\mathbf{v} | \mathbf{u}) P_{\mathbf{W} | \mathbf{U}}(\mathbf{w} | \mathbf{u}). \quad (2.6)$$

See Figure 2.3 for a visualization of a graph that satisfies (2.6).

The global Markov property entails strong restrictions on the factorization properties of the joint distribution. Specifically, according to the *Hammersley-Clifford-Theorem* [HC71], the joint distribution factorizes into a set of *potentials* $\Psi = \{\Phi_{C_1}, \dots, \Phi_{C_K}\}$ specified over the set of cliques $\mathbf{C} = \{C_1, \dots, C_K\}$ of the graph. That is

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \prod_{C_i \in \mathbf{C}} \Phi_{C_i}(\mathbf{x}_{C_i}), \quad (2.7)$$

where the potentials $\Phi_{C_i}(\mathbf{x}_{C_i})$ only depend on the values of random variables belonging to the given clique, i.e., $\mathbf{x}_{C_i} = \{x_i : X_i \in C_i\}$. Although the potentials need to be non-negative and may remind us of conditional probabilities, it is important to stress that this is not the case and that potentials do not necessarily have a specific probabilistic interpretation (cf. [KF09, Example 4.2]). The flexibility one has in assigning values to the potential function comes at a cost, however: the product over all potential functions may require a normalization function (partition function) \mathcal{Z} to ensure that the product in (2.7) constitutes a valid distribution.⁵

The product in (2.7) can, in principle, consist of potentials specified only over the max-cliques of the graphs. This may, however, hide the factorization properties of the underlying problem [KF09, p.108] and it is thus often preferable to consider only smaller (sub)-cliques as

⁴ We consider only random variables with a discrete alphabet \mathcal{X} although the framework would generalize to continuous random variables with $\mathcal{X} = \mathbb{R}$ as well.

⁵ The partition function is evaluated by computing the unnormalized sum over all states and is usually denoted by the letter \mathcal{Z} to denote its origin in the German word *Zustandssumme*.




Figure 2.3: Graphical representation of the conditional independencies implied by (2.6). The set of nodes \mathbf{U} separates \mathbf{V} from \mathbf{W} , i.e., the after removing the subgraph induced by \mathbf{U} from \mathcal{G} two separate subgraphs remain.

e.g., of size two for *pairwise models*. Pairwise models are particularly well-suited for a theoretical analysis because of their simplicity and will therefore be the main focus of this thesis. Because of their central role, we devote the following section to the introduction of binary pairwise models.

2.4 Binary Pairwise Models

Throughout the main part of this thesis, we will focus on one particular type of undirected models: these are *binary pairwise graphical models*. Binary pairwise models admit a relatively simple treatment while being rich enough to represent a wide class of problems. Indeed, most practical problems permit a representation as a binary pairwise graphical models, which highlights that the class is much less restrictive as it may seem at first.

We will first provide a formal introduction in Section 2.4.1 and then discuss both the limitations and the relevance of binary pairwise graphical models in Section 2.4.2. Finally, we introduce the exponential representation in Section 2.4.3 and clarify the relation to the Ising model in Section 2.4.4.

2.4.1 Model Description

Binary models are graphical models where every random variable $X_i \in \mathbf{X}$ takes values from the binary alphabet, e.g., according to $x_i \in \mathcal{X} = \{-1, +1\}$. *Pairwise* models have potential functions that consist of two random variables at most; i.e., potentials are only associated with cliques $C_i \in \mathbf{C}$ of size $|\mathbf{X}_{C_i}| \leq 2$. We resort to an even finer-grained representation that separately specifies the potentials over the nodes and edges. Complying to this representation, we then

define the joint distribution according to

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \prod_{(i,j) \in \mathbf{E}} \Phi_{X_i, X_j}(x_i, x_j) \prod_{X_i \in \mathbf{X}} \Phi_{X_i}(x_i), \quad (2.8)$$

where we introduced two different potential-types. These are: the pairwise potentials $\Phi_{X_i, X_j}(x_i, x_j)$ that act on the edges $(i, j) \in \mathbf{E}$ and the local potentials $\Phi_{X_i}(x_i)$ that act on the nodes $X_i \in \mathbf{X}$. Note that each pairwise potential is only considered once as $(i, j) = (j, i)$. We will prefer to denote the local and pairwise potentials by their shorthand notation $\Phi(x_i)$ and $\Phi(x_i, x_j)$, unless the actual values of the associated random variables are of immediate relevance.

Note that the factorization in (2.7) considers the set of all cliques (with varying sizes), whereas the factorization in (2.8) considers cliques of size two at most. Therefore, note that every connected pair of variables is also a clique (albeit possibly only a subset of a larger clique). This means that both forms become identical – even in the presence of larger cliques – whenever all higher-order potentials are trivial.⁶

The representation in (2.8) has yet another equivalent factorization

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \prod_{C_{(i,j)} \in \mathbf{C}} \Phi_{C_{(i,j)}}(x_i, x_j)$$

that incorporates all local potentials into the pairwise ones. Such a compact representation clearly reduces the overall number of potentials, but may obscure the underlying structure. This often makes the model much less intuitive to interpret. We will now discuss and compare the different ways of factorizing the joint distribution by means of the following example.

Example 1 (Factorization into Potentials of Varying Size).

Consider the undirected graphical model $\mathcal{U} = (\mathcal{G}, \Psi)$ depicted in Figure 2.4.




Figure 2.4: Complete graph with three random variables with its potentials specified over the set of edges and the set of nodes according to (2.8).

We consider a joint distribution that is specified according to (2.8) in terms of local potentials $\Phi(x_i)$ and pairwise potentials $\Phi(x_i, x_j)$ so that

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \Phi(x_1, x_2) \Phi(x_2, x_3) \Phi(x_3, x_1) \Phi(x_1) \Phi(x_2) \Phi(x_3). \quad (2.9)$$

As discussed above, one possible factorization stems from incorporating the local potentials into the pairwise ones. The specific model in Figure 2.4 provides one obvious way of doing

⁶ A trivial potential does not influence the values of the joint distribution. One way of making a potential trivial is by assigning identical values, as for example $\Phi_C(\mathbf{x}_C) = 1$, for all possible configurations \mathbf{x}_C .

so; therefore, let us define the clique potentials according to

$$\Phi_{C_{(i,j)}}(x_i, x_j) = \Phi(x_i, x_j)\Phi(x_j) \quad (2.10)$$

so that

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \Phi_{C_{(1,2)}}(x_1, x_2)\Phi_{C_{(2,3)}}(x_2, x_3)\Phi_{C_{(3,1)}}(x_3, x_1). \quad (2.11)$$

Note that (2.10) is just one possible way of constructing clique potentials from the local and the pairwise ones. One only has to make sure that every potential is considered exactly once, as the joint distribution in (2.11) would not conform with its original definition in (2.9) otherwise.

Finally, it is always possible to factorize the joint distribution over the maximum cliques. Note that, in this example, \mathcal{G} is already a maximum clique, which makes it possible to express the joint distribution by a single clique potential. This potential of the maximum clique incorporates all pairwise- and singleton- potentials according to

$$\Phi_{C_{(1,2,3)}}(x_1, x_2, x_3) = \Phi(x_1, x_2)\Phi(x_2, x_3)\Phi(x_3, x_1)\Phi(x_1)\Phi(x_2)\Phi(x_3) \quad (2.12)$$

so that

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\mathcal{Z}} \Phi_{C_{(1,2,3)}}(x_1, x_2, x_3). \quad (2.13)$$

The above example highlights the influence of the particular factorization on the overall number of potentials. Higher-order cliques reduce the overall number of potentials but, simultaneously, tend to conceal the underlying structure of the joint distribution [KF09, Chapter 4.2]). Similarly, pairwise graphical models enforce one specific factorization and may thus hide the “true” factorization of the underlying distribution as well. Factor graphs, on the other hand, provide a flexible representation that makes this underlying factorization explicit. We will, however, except for the application to error-correcting codes, restrict ourselves to pairwise graphical models for the remainder of this thesis.

2.4.2 Generality of Binary Pairwise Models

The simplicity of binary pairwise models is compelling, although the exclusion of other models – not representable by binary pairwise models – may seem rather restrictive. Despite this impression, binary pairwise models provide a general framework. That is, most graphical models can be directly converted into a binary pairwise model [YFW01, Wei00].

The conversion from general undirected models to binary pairwise models becomes immediately apparent for models with strictly positive potentials, in which case, a simple reduction to the binary pairwise case exists [EG13]. Note, however, that such a conversion potentially scales up the problem immensely and is therefore not always ideal from a practical point of view [Mac01].

From a theoretical point of view, however, it is beneficial to study binary pairwise models; these models are rich enough to exhibit complex behavior – thus demonstrating many interesting aspects, though they still admit a simplified treatment – thus avoiding many technical subtleties. Moreover, the study of binary pairwise models carries great relevance; not only because they are still far from being fully understood but also because theoretical insights carry over to more general models (by conversion of one model-class to the other).

Besides their generality, binary pairwise models are important in their own right as they arise in various applications.

2.4.3 Exponential Representation and Model Parametrization

So far we have not considered the actual specification of the potentials. We will now introduce one specific parameterization of binary pairwise models, show why it belongs to an exponential family, and define different model-classes with fundamentally different behavior. Besides, we will fix our naming convention and briefly point to alternative ones. Note that, because of its immediate connection, the terminology is often rooted in the physical interpretations of the quantities.

Many relevant problems have a non-zero probability for all configurations \mathbf{x} , such that the joint distribution factorizes into strictly positive potentials. This allows us to assign some *energy* $E(\mathbf{x})$ to every configuration, so that the joint distribution from (2.8) can be expressed in its exponential form according to

$$P_{\mathbf{X}}(\mathbf{x}) = \frac{1}{Z} \cdot e^{-E(\mathbf{x})}. \quad (2.14)$$

For the remainder of this thesis we shall use a minimal representation of binary pairwise models (cf. [WJ08, Section 3.3]) that defines the energy in terms of couplings $J_{ij} \in \mathbb{R}$, that act on the edges $(i, j) \in \mathbf{E}$ and local fields $\theta_i \in \mathbb{R}$, that act on the nodes $X_i \in \mathbf{X}$. Note that we drop the subscripts and write $J_{ij} = J$ or $\theta_i = \theta$ whenever the parameters are identical for all edges or for all nodes respectively.⁷

Let the local and pairwise potentials of state $x_i \in \{-1, +1\}$ be $\Phi(x_i) = \exp(\theta_i x_i)$ and $\Phi(x_i, x_j) = \exp(J_{ij} x_i x_j)$. Then, if we plug these potentials into (2.8) we end up with a joint distribution that has its energy (cf. (2.14)) given by

$$E(\mathbf{x}) = - \sum_{(i,j) \in \mathbf{E}} J_{ij} x_i x_j - \sum_{X_i \in \mathbf{X}} \theta_i x_i. \quad (2.15)$$

The energy (and consequently the joint distribution) depends not only on \mathbf{x} but on J_{ij} and θ_i as well; we will, however, only make this dependence explicit by $E(\mathbf{x}, J, \theta)$ if of immediate relevance.

In the literature, one distinguishes two different types of interactions between variables: if a coupling is positive ($J_{ij} > 0$) then the associated edge (i, j) is *attractive*; if a coupling is negative ($J_{ij} < 0$) then the associated edge (i, j) is *repulsive*. In accordance with this naming-convention, we call a model \mathcal{U} attractive if it contains only attractive edges (these models are also known as ferromagnetic models [MM09] or log-supermodular models [Ruo12]); we call it repulsive (or antiferromagnetic) if it contains only repulsive edges; and we call a model *general* if it contains both types of edges.

The particular representation of (2.15) is also known as Hopfield network (cf. [Hop82] and [Mac03, Chapter 42]) or Boltzmann machine (cf. [HS86, WT03] and [Mac03, Chapter 43]) in the machine learning community and as Ising model (cf. [Bru67] and [Mac03, Chapter 31]) in the physics literature. The Ising model has been studied for a long time in statistical physics. For its particular relevance, we will devote the subsequent section to the Ising model, discuss its most important properties, and provide a brief historical outline of its development.

⁷ If all nodes have the same value $\theta_i = \theta$ we will sometimes refer to θ as the external field.

2.4.4 The Ising Model

The study of pairwise models in the form of (2.14) actually dates back for more than a century.⁸ In the field of statistical physics, one studies the macroscopic behavior of systems with a large number of interacting components (e.g., atoms or molecules) in the thermodynamic limit of infinitely many components.

One concept of central relevance are phase transitions; these are points in the parameter space where the partition function becomes non-analytic. Phase transitions trigger fundamental changes in the system behavior, as for example the change from a fluid to gas. In computer science, we often encounter algorithms that have many “components” interacting with each other (e.g., as in message passing algorithms); such algorithms often exhibit rapid performance-drops for certain points in the parameter space, which clearly remind us of phase transitions.

One extensively studied model in statistical physics is the Ising model: it consists of N atoms in a configuration $\sigma = (\sigma_1, \dots, \sigma_N)$ with spins $\sigma_i \in \{-1, +1\}$ that lie on a d -dimensional lattice and that are magnetically coupled.⁹ As a model for magnetic bodies, the Ising model’s phase transitions describe the change from the paramagnetic to the ferromagnetic region among others. The Ising model is one of the simplest models that is sufficiently rich to exhibit phase transitions; therein lies its appeal.

The interaction between two neighboring spins is specified by the coupling strength J that is either ferromagnetic (i.e., positive) or anti-ferromagnetic (i.e., negative). Two neighbors then energetically favor the same state in case of ferromagnetic interactions and the opposite state in case of anti-ferromagnetic interactions. The energy of a configuration σ is

$$E(\sigma) = -\beta J \sum_{(i,j) \in \mathbf{E}} \sigma_i \sigma_j + H \sum_{i=1}^N \sigma_i, \quad (2.16)$$

where H is the external magnetic field and $\beta = \frac{1}{k_B T}$ is the inverse temperature with k_B being the Boltzmann constant.

The appreciation of the Ising model had its ups and downs. After being dumped as a too simplified model with no physical usefulness, it took some years before the Ising model received its well-deserved attention. We outline some of the cornerstones in this development below and refer the interested reader to the surprisingly exciting read [Bru67] for a review of the major events in the history of the Ising model.

Wilhelm Lenz proposed the Ising model, as a simplified model for the interactions inside magnetic bodies, with the aim of developing a better understanding of the underlying properties. One of his students, Ernst Ising, solved the one-dimensional case (by applying the transfer-matrix method) in his thesis. The result, rather surprisingly, revealed the inexistence of phase-transitions in the one-dimensional case. Being unable to solve it in higher dimensions, he concluded that the Ising model is generally incapable of experiencing phase transitions. This conjecture, however, was in stark contrast to the ferromagnetic theory developed by Pierre Currie, which led to the disregard of the Ising model.

Finally, this contradiction was resolved by Lars Onsanger who solved the Ising model on the two-dimensional grid in a mathematical “tour-de-force” (as the Onsanger solution is nowadays often referred to) and revealed the existence of phase transitions in two dimensions. Note that the exact solution of the Ising model is only known in these two cases and it remains an open

⁸ Distributions of this form were first introduced by Ludwig Boltzmann and by Josiah Willard Gibbs and provide the foundation of statistical physics. In the physics literature, one often refers to distributions of the form (2.14) as Boltzmann-distributions or Gibbs-measures.

⁹ Here we use the standard notation for the Ising model, although the correspondence to the notation in this thesis becomes obvious by comparing (2.15) to (2.16).

problem to compute the exact solution for higher-dimensional cases.

The classical Ising model, with its energy defined according to (2.16), provides a powerful generalization: the *spin glass* model has the (local) fields θ_i and the couplings J_{ij} take potentially different values for all nodes and edges and has its energy defined according to

$$E(\sigma) = -\beta \sum_{(i,j) \in \mathbf{E}} J_{ij} \sigma_i \sigma_j + \sum_{i=1}^N \theta_i \sigma_i. \quad (2.17)$$

What makes these spin glass models so interesting is that we have a much poorer understanding of them as opposed to the classical Ising model. The lack of understanding is mainly because of *frustrations* [Gé77]. These are configurations that have some pairs of random variables X_i, X_j energetically favor states that contradict the coupling of the associated edge (i, j) (cf. Example 4). To make this more precise we introduce the product of couplings along some cycle \mathcal{P} , i.e.,

$$J_{\mathcal{P}} = \prod_{\mathcal{P}} J_{ij}. \quad (2.18)$$

Then, a graph is frustrated whenever it contains a cycle for which $J_{\mathcal{P}}$ equates to a negative number. For the classical Ising model where all couplings take the same value, the existence of frustrations depends only on the graph structure (i.e., if cycles of odd-length exist). For spin glasses, however, the existence of frustrations depends on both the graph structure and the parameters. This is one of the main reasons for studying spin glasses as frustrations can lead to a complex solution space with potentially many different solutions.

So far, we have seen that binary pairwise models constitute a very general class of graphical models and fit nicely into the framework of statistical physics. Although this emphasizes the relevance of binary pairwise models, we barely mentioned the purpose of introducing graphical models. We will now shift our focus to the problems we intend to solve and thus reveal the elegance and the advantages of graphical models.

2.5 Inference

The task of inference plays an important role in many scientific fields and it is a prerequisite for probabilistic reasoning. In a nutshell, inference deals with drawing statistical conclusions about a certain subset of random variables, given some (possibly noisy) observations [PPT14].

Applications include, but are not limited to, computer vision and speech processing where the observations are corrupted versions of the image or the speech-signal and one is interested in finding the most probable explanation. Ideally, one would hope that the best explanation for the given observations matches the original image or signal. Another important application is found in the context of error-correcting codes, where the observation is the received codeword and one is interested in obtaining the sent codeword.

To make the task of inference more precise we consider a joint distribution $P_{\mathbf{X}}(\mathbf{x})$ over a set of nodes \mathbf{X} that consists of two disjoint subsets, i.e, $\mathbf{X} = \mathbf{Y} \cup \mathbf{O}$ and $\mathbf{Y} \cap \mathbf{O} = \emptyset$. The observed variables are denoted by \mathbf{O} and the unobserved ones are denoted by \mathbf{Y} . Without loss of generality, we will restrict ourselves to problems without observed variables $\mathbf{O} = \emptyset$ throughout this thesis. Now we present the following three central problems of probabilistic inference.

- *Maximum a posteriori (MAP) inference* estimates the mode of the distribution with the aim of identifying the joint assignment \mathbf{y}^* that maximizes the probability according to

$\mathbf{y}^* = \arg \max_{\mathbf{y} \in \mathcal{Y}^{|\mathbf{Y}|}} P_{\mathbf{Y}|\mathbf{O}}(\mathbf{y}|\mathbf{o})$. Note that, strictly speaking, this is the most probable explanation (MPE) estimate, whereas the MAP estimate deals with the more general problem of maximizing the probability for any subset of \mathbf{Y} ; it is quite common, however, to neglect this distinction despite MPE inference being the easiest instance [KFD⁺07, Chapter 2.3]. We abide to this convention and, provided $\mathbf{O} = \emptyset$, aim to obtain

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}^N} P_{\mathbf{X}}(\mathbf{x}). \quad (2.19)$$

- *Marginal inference* is the task of computing the marginal distribution for a subset of random variables $\mathbf{Y} \subset \mathbf{X}$, i.e., to compute $P_{\mathbf{Y}}(\mathbf{y})$ according to (2.2). Note that we will sometimes use the following shorthand notation for this double-summation where

$$P_{\mathbf{Y}}(\mathbf{y}) = \sum_{\mathbf{X} \setminus \mathbf{Y}} P_{\mathbf{X}}(\mathbf{x}) = \sum_{\mathbf{O}} P_{\mathbf{X}}(\mathbf{x}) = \sum_{X_i \in \mathbf{O}} \sum_{x_i \in \mathcal{X}} P_{\mathbf{X}}(\mathbf{x}). \quad (2.20)$$

We will be particularly interested in the singleton marginals $P_{X_i}(x_i)$ for single random variables and the pairwise marginals $P_{X_i, X_j}(x_i, x_j)$ for pairs of random variables.

For binary models it is often more convenient to work with the expectations, which are the mean m_i (or magnetization) and the correlation χ_{ij} , instead of considering the singleton marginals $P_{X_i}(x_i)$ and the pairwise marginals $P_{X_i, X_j}(x_i, x_j)$ explicitly, where

$$m_i = \mathbb{E}(X_i) = P_{X_i}(X_i = 1) - P_{X_i}(X_i = -1), \quad (2.21)$$

$$\chi_{ij} = \mathbb{E}(X_i X_j). \quad (2.22)$$

- Another important problem is to evaluate the *partition function*

$$\mathcal{Z} = \sum_{\mathbf{x} \in \mathcal{X}^N} \prod_{C_i \in \mathbf{C}} \Phi_{C_i}(\mathbf{x}_{C_i}), \quad (2.23)$$

which is the normalization coefficient of the joint distribution in (2.7).

We will primarily focus on the problems of computing marginal distributions and evaluating the partition function in this thesis. Note that these two problems are in fact closely related as the marginal distribution $P_{X_i}(x_i)$ equals the ratio between a partial partition function $\mathcal{Z}(X_i)$ and the partition function:

$$\begin{aligned} P_{X_i}(x_i) &= \sum_{\mathbf{X} \setminus X_i} P_{\mathbf{X}}(\mathbf{x}) = \sum_{\mathbf{X} \setminus X_i} \frac{1}{\mathcal{Z}} \prod_{C_i \in \mathbf{C}} \Phi_{C_i}(\mathbf{x}_{C_i}), \\ &= \frac{1}{\mathcal{Z}} \sum_{\mathbf{X} \setminus X_i} \prod_{C_i \in \mathbf{C}} \Phi_{C_i}(\mathbf{x}_{C_i}), \\ &= \frac{\sum_{\mathbf{X} \setminus X_i} \prod_{C_i \in \mathbf{C}} \Phi_{C_i}(\mathbf{x}_{C_i})}{\sum_{\mathbf{x} \in \mathcal{X}^N} \prod_{C_i \in \mathbf{C}} \Phi_{C_i}(\mathbf{x}_{C_i})} = \frac{\mathcal{Z}(X_i)}{\mathcal{Z}}. \end{aligned} \quad (2.24)$$

The computation of the marginals seems relatively straightforward according to (2.24). Yet, there is one fundamental problem that prohibits the application of (2.24) to practical problems, which essentially boils down to the overall number of variables involved. Even if the joint distribution $P_{\mathbf{X}}(\mathbf{x})$ is known – neglecting the fact that the memory complexity of storing $P_{\mathbf{X}}(\mathbf{x})$ is exponential in the number of variables – computing $P_{X_i}(x_i)$ is problematic. More specifically,

consider a joint distribution $P_{\mathbf{X}}(\mathbf{x})$ specified over N random variables with $k = |\mathcal{X}|$ states. Then, the sum in the numerator of (2.24) is evaluated $N - 1$ times and goes over k terms each time; i.e., in total computing a marginal distribution would require the summation over k^{N-1} terms in total.

This drastically limits the problem size for which the marginals can be evaluated in practice. On the positive side, however, we have already encountered the compact representation of probabilistic graphical models. We will subsequently show how to utilize this representation and how this opens the door for efficient inference methods. Before considering arbitrary models, we focus on models with certain graph structures that lend themselves to particularly elegant ways of performing inference. Ideally one should aim to exploit this representation beyond that and extend it to the central problems of inference.

2.5.1 Exact Inference: Efficient Methods

Tree-Structured Graphs

Chains and trees possess Markov properties that, if exploited properly, give rise to efficient inference methods. The statistical dependencies imposed by tree-structured models admit inference methods that, instead of manipulating the joint distribution directly, recursively perform local computations; this reduces the computational complexity immensely. Various inference methods were independently introduced in different fields (cf. Section 3.1 for a brief overview) that all rely on the very same principle to perform efficient inference. The basic principle is to perform a set of local computations, often interpreted as messages between random variables; this also explains the term *message passing algorithms* that is often used to unite all those algorithms. The following example demonstrates the underlying principles and highlights the efficiency of message passing algorithms.

Example 2 (Exact Inference on a Tree).

Consider the undirected graphical model $\mathcal{U} = (\mathcal{G}, \Psi)$ depicted in Figure 2.5. First, we express the joint distribution as the product over all pairwise clique potentials according to (2.7) so that

$$\begin{aligned} P_{\mathbf{X}}(\mathbf{x}) &= \frac{1}{Z} \prod_{C_i \in \mathcal{C}} \Phi_{C_i}(\mathbf{x}_{C_i}) \\ &= \frac{1}{Z} \Phi_{C_{(1,4)}}(x_1, x_4) \Phi_{C_{(2,4)}}(x_2, x_4) \Phi_{C_{(3,4)}}(x_3, x_4) \Phi_{C_{(4,5)}}(x_4, x_5) \cdot \\ &\quad \Phi_{C_{(5,6)}}(x_5, x_6) \Phi_{C_{(6,7)}}(x_6, x_7). \end{aligned} \quad (2.25)$$

Note that the factorization in terms of pairwise cliques equals the factorization in terms of maximum cliques for tree-structured models for the lack of loops.

Second, we compute the singleton marginals for one specific random variable, e.g., X_5 , by summing over all other variables such that

$$P_{X_5}(x_5) = \frac{1}{Z} \sum_{x_1 \in \mathcal{X}} \cdots \sum_{x_4 \in \mathcal{X}} \sum_{x_6 \in \mathcal{X}} \sum_{x_7 \in \mathcal{X}} \cdot \prod_{C_i \in \mathcal{C}} \Phi_{C_i}(\mathbf{x}_{C_i}) \quad (2.26)$$

A closer look at the clique potentials reveals the benefits of reordering of the summations. The total amount of summations becomes much more manageable if we make use of the commutative and the distributive law and rewrite (2.26) according to

$$\begin{aligned}
P_{X_5}(x_5) = & \left(\sum_{x_4 \in \mathcal{X}} \Phi_{C(4,5)}(x_4, x_5) \cdot \right. \\
& \underbrace{\left(\sum_{x_1 \in \mathcal{X}} \Phi_{C(1,4)}(x_1, x_4) \right)}_{\mu_{14}(x_4)} \underbrace{\left(\sum_{x_2 \in \mathcal{X}} \Phi_{C(2,4)}(x_2, x_4) \right)}_{\mu_{24}(x_4)} \underbrace{\left(\sum_{x_3 \in \mathcal{X}} \Phi_{C(3,4)}(x_3, x_4) \right)}_{\mu_{34}(x_4)} \Big) \\
& \underbrace{\hspace{10em}}_{\mu_{45}(x_5)} \\
& \left. \underbrace{\left(\sum_{x_6 \in \mathcal{X}} \Phi_{C(5,6)}(x_5, x_6) \right)}_{\mu_{65}(x_5)} \underbrace{\left(\sum_{x_7 \in \mathcal{X}} \Phi_{C(6,7)}(x_6, x_7) \right)}_{\mu_{76}(x_6)} \right) . \tag{2.27}
\end{aligned}$$

Note that we introduced the powerful notion of messages in (2.27); messages $\mu_{ij}(x_j)$ provide a compact notational convention and express the consequence of summing over the respective variable. We say that $\mu_{ij}(x_j)$ is passed along the edge from X_i to X_j (see Figure 2.5).

The desired marginal distribution $P_{X_5}(x_5)$ is then given by the normalized product of all incoming messages. We consequently realize from (2.27) that

$$P_{X_5}(x_5) = \frac{1}{Z} \mu_{45}(x_5) \mu_{65}(x_5). \tag{2.28}$$




Figure 2.5: Tree-structured model that illustrates the concept of message passing with the purpose of computing the marginal distribution $P_{X_5}(x_5)$. All messages of (2.27) are illustrated next to the associated edges.

Note how the messages $\mu_{45}(x_5)$ and $\mu_{65}(x_5)$ incorporate all messages of the sub-trees rooted in X_4 and X_6 . Consequently, knowledge of only the two messages $\mu_{45}(x_5)$ and $\mu_{65}(x_5)$ is sufficient to compute the marginal $P_{X_5}(x_5)$ according to (2.28).

The reordering of the summations reduces the overall amount of required summations notably. Comparison of (2.26) and (2.27) reveals that the computational complexity reduces from $\mathcal{O}(|\mathcal{X}|^N)$ to $\mathcal{O}(|\mathcal{X}|^2)$.

Remember how $\mu_{45}(x_5)$ accounted for all messages that come from the subgraph induced by $\mathbf{X}' = \{X_1, X_2, X_3, X_4\}$. This suggests a recursive rule for computing the messages without the need for explicitly rearranging the summations. In particular, we can compute the message $\mu_{ij}(x_j)$ by taking the product of all incoming messages, except the one from X_j , times the

pairwise potential $\Phi_{C(i,j)}(x_i, x_j)$ so that

$$\mu_{ij}(x_j) = \sum_{x_i \in \mathcal{X}} \Phi_{C(i,j)}(x_i, x_j) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}(x_i). \quad (2.29)$$

The message $\mu_{ij}(x_j)$ is a vector over all states $x_j \in \mathcal{X}$. One can interpret this message as the belief of node X_i about the relative probabilities that X_j is in state x_j given all the information¹⁰ that is available to X_i , except from X_j .

The notion of messages has a further advantage: if we wish to compute the marginals $P_{X_i}(x_i)$ for multiple nodes in the graphical model we can, instead of performing the procedure multiple times, reuse the already computed messages. It suffices to pass the messages back and forth throughout the whole graph only once. For a chain this means propagating the messages along both directions.¹¹ If we wish to apply this procedure to a tree we must first fix some ordering, i.e., pick one node as root, before then passing the messages upwards from the leaves and then pass the messages down again.

The representation of the summations in (2.27) in terms of messages captures an important property of tree-structured graphical models. For every $X_i \in \mathbf{X}$ it holds that any pair of its neighbors $\{X_k, X_l\} \in \partial(i)$ is conditionally independent given X_i , i.e.,

$$P_{X_k, X_l | X_i}(x_k, x_l | x_i) = P_{X_k | X_i}(x_k | x_i) P_{X_l | X_i}(x_l | x_i). \quad (2.30)$$

These statistical independence statements further impose one important property on the graph. Namely, that removing the edge (i, k) creates two disconnected subgraphs $\mathcal{G}^{(i)} = (\mathbf{X}^{(i)}, \mathbf{E}^{(i)}) : X_i \in \mathbf{X}^{(i)}$ and $\mathcal{G}^{(k)} = (\mathbf{X}^{(k)}, \mathbf{E}^{(k)}) : X_k \in \mathbf{X}^{(k)}$. Note that the existence of any alternative path between $\mathcal{G}^{(i)}$ and $\mathcal{G}^{(k)}$, i.e., the existence of loops, would violate (2.30).

Graphs with Loops

So far we have witnessed the capabilities of message passing algorithms for chains and tree-structured models; one might ask how this concept generalizes and if it is able to cope with loops. This is a critical question for two reasons: efficient inference methods are essential when confronted with large and loopy models; and many problems of practical relevance correspond to loopy models. Every general-purpose inference method must, therefore, be able to cope with loops.

Unfortunately, this is not the case for message passing algorithms. Loopy graphs violate the independence statements in (2.30) and thus prevent a straightforward generalization of message passing to loopy graphs.

As discussed, excluding models with loops is not a viable option either. Nonetheless, exploiting the Markov properties seemed promising. Fortunately, we can have both. That is, an algorithm exists that is both capable of accounting for loops and capable of exploiting the Markov properties. These seemingly conflicting intentions are satisfied by modifying the graph and carefully constructing certain subgraphs until we finally end up with a tree again.

The failure of message passing is demonstrated on a loopy graph in Example 3. Subsequently, we show how the graph needs to be modified to allow for efficient inference.

¹⁰ We are using the term information here only figuratively in the sense that all necessary summations for evaluating $P_{X_i}(x_i)$ are subsumed in $\mu_{ji}(x_i)$. This does not adhere to the formal definition of entropy as an information-theoretic measure.

¹¹ In the particular case of a chain the messages actually equal the Chapman Kolmogorov equations.

Example 3 (Exact Inference on Loopy Graphs).

We create a loopy graph by adding the edge $(5, 7)$ to the model of Example 2 (cf. Figure 2.6 (a)). Let us focus on the subgraph induced by $\mathbf{X}' = \{X_5, X_6, X_7\}$. We first observe that the conditional independence statements of (2.30) are violated as

$$P_{X_5, X_7 | X_6}(x_5, x_7 | x_6) \neq P_{X_5 | X_7}(x_5 | x_7) P_{X_6 | X_7}(x_6 | x_7); \quad (2.31)$$

and second that information flows from X_7 towards X_5 (and the remainder of the graph) via two different paths: once via $\mathcal{P}_1 = (\{X_5, X_7\}, \{(7, 5)\})$ and once via $\mathcal{P}_2 = (\{X_5, X_6, X_7\}, \{(7, 6), (6, 5)\})$. As a result, the message $\mu_{57}(x_7)$, if computed according to (2.29), incorporates the belief stemming from X_7 via one path and erroneously propagates it back to X_7 via the other path.




Figure 2.6: (a) Model from Example 2 with the additional edge $(5, 7)$ that creates a loop. (b) Junction tree for the same model; note that variable grouping renders the graph tree-structured again.

The graphical model needs to be tree-structured in order to satisfy the formal requirements for message passing to work; this, however, is not the case as seen in Figure 2.6(a). We can, however, form super-nodes, consisting of multiple variables, until no more loops are present in the graph to leverage the power of message passing. For our current example one can simply group the variables X_5, X_6 , and X_7 together into X_{5-6-7} and define the pairwise potential between X_4 and this super-node according to

$$\Phi_{C_{(4,5-6-7)}}(x_4, x_5, x_6, x_7) = \Phi_{C_{(4,5)}}(x_4, x_5) \Phi_{C_{(5,6)}}(x_5, x_6) \Phi_{C_{(6,7)}}(x_6, x_7) \Phi_{C_{(5,7)}}(x_5, x_7). \quad (2.32)$$

The resulting graph is depicted in Figure 2.6 (b). Now that we have created a tree-structured model again (cf. Figure 2.6(b)), we can simply perform message passing on this modified graph and compute the marginals accordingly. Note how the new node X_{5-6-7} has an increased number of states according to $|\mathcal{X}|^{|C_{5-6-7}|} = |\mathcal{X}|^3$ now. Despite the exponential growth of the state space, grouping variables together has its merits. In particular, it suggests an optimal way of rearranging the variables in the computation of the marginals that maintains the Markov properties despite the existence of loops.

We have seen in Example 3 that performing inference on loopy graphs increases the computational complexity. This is an immediate consequence of considering higher-order cliques in the construction of the *junction tree* (i.e., the modified graph). In general – and for models with multiple intertwined loops in particular – however, there is not just one but many possibilities of grouping variables together. Consequently, one should construct the junction tree such that the largest clique is kept as small as possible. If one wants to keep the complexity in check, it is thus of utmost importance to carefully construct the junction.

One might ask whether a modified graph (i.e., a clique- or junction-tree) without loops always

exists – and how it is constructed. Indeed, for every graphical model such a junction tree exists and the construction of an optimal junction tree – with as small cliques as possible – involves three major steps. These steps define the *junction tree algorithm* [LS88] that we will review below. We do not give an exhaustive introduction with all the details but rather aim to provide an overview that highlights the overall idea of the required operations.

First, note that every graph \mathcal{G} has an associated junction tree if and only if \mathcal{G} is triangulated [LS88]. A graph is triangulated (or chordal) if every cycle of length four or greater has a chord (i.e., an edge that joins two nodes of the cycle). Usually, a graph possesses multiple valid triangulated graphs, the choice of which has a major influence on the overall computational complexity [KF09, Section 10.4]. Finding an optimal, that is a minimal, triangulation is NP-hard on its own but various heuristics exist that find reasonably good triangulation. Most of these heuristics rely on a particular form of variable elimination, whereas it depends on the given graph which method performs best [KF09, Section 9.4.3.2].

Second, the triangulated graph provides the basis for the construction of the junction tree. The nodes in the junction tree correspond to maximal cliques in the triangulated graph; if a variable is present in two cliques, the cliques are joined by an edge in the junction tree. A triangulated graph usually admits multiple junction trees of varying sizes. The construction of the junction tree thus has a major influence on the overall efficiency. Typically, one assigns a weight to each edge that corresponds to the number of variables included in both cliques and subsequently constructs a maximum spanning tree [KF09, Section 10.4.2].

Finally, once the junction tree is constructed an efficient inference method, for example message passing, can be applied to the junction tree to yield the marginal distributions over all cliques. The singleton marginals are then obtained by direct summation over a clique containing the desired variable.

The junction tree algorithm is straightforward in principle and provides a graph-based approach that exploits the factorization properties for loopy graphs. Although the junction tree significantly reduces the complexity of exact inference, it is still of limited practical use. This stems from working with the maximal cliques of the triangulated graph and the exponential growth of the state space with the clique-size. The applicability of the junction tree algorithm is consequently limited by the size of its largest clique [BKvdEvdG01].

Nonetheless, there are two reasons that justify the introduction of the junction tree: First, despite its limitation to problems with relatively small tree-width, the junction tree is efficient in some sense; in particular, no general exact inference method exists that is computationally more efficient than the junction tree [Bis06, Section 8.4.6]. The junction tree thus serves as the method of choice for estimating the ground truth when assessing approximate inference methods in the subsequent chapters. Second, the junction tree algorithm reveals how the inherent properties of loopy graphs increase the complexity of exact inference; this further emphasizes the need for efficient approximation methods.

In the subsequent chapter, we will explicitly focus on loopy graphs and show how to tackle the related complexity issues. We take different points of view on performing approximate inference, extend the underlying concept of message passing to more general graphs, and discuss how these methods cope with the existence of loops. Ultimately, this search for efficient methods culminates in a wide range of available approximate inference methods.

3

Approximate Inference: Belief Propagation

*“The stars bend like slaves to laws not
decreed for them by human intelligence,
but gleaned from them.”*
– Ludwig E. Boltzmann

Owing to belief propagation’s specific relevance for this thesis, we present belief propagation (BP) in detail. This includes the classical definition as a message passing algorithm in Section 3.2, as well as the variational interpretation in Section 3.3 that highlights the connection to methods from statistical physics. Moreover, we cast BP as a dynamical system and elaborate on the resulting implications in Section 3.5. The perspective of dynamical systems gives rise to a whole group of message passing algorithms for approximate inference, where every particular instance comes with its own pitfalls and limitations. We briefly discuss some popular variants of BP and describe how each one enhances the performance over the standard implementation in Section 3.5.2- 3.5.4.

Much of this chapter summarizes established textbook knowledge (cf. [KF09, MWJ99, MM09]) although we review some of the most recent developments as well. The representation of BP as a dynamical system is rather obvious and has thus been considered multiple times (e.g., in [MK05, RKDW10, TR06]); here, we emphasize the added value of doing so and show how casting BP as a dynamical system provides a unifying framework. The presentation of BP variants in Section 3.5.3 contains results developed in collaboration with Michael Rath, Sebastian Tschatschek, and Franz Pernkopf [KRTP15].

3.1 Motivation

In Chapter 2 we have seen how exact inference methods suffer from the existence of loops. In fact, exact inference is NP-hard [Coo90], unless the probabilistic graphical model is appropriately restricted, which, however, would rule out many models of practical relevance. This highlights the need for efficient approximate inference methods. Even approximate inference, however, is NP-hard if a certain accuracy is required [Rot96, DL93]. Note, however, that this pessimistic statement does not render approximate inference completely useless, but only indicates that specific models do exist for which achieving the desired accuracy is intractable.

If one wants to perform approximate inference, one can choose from a wide range of different methods. Before one can make a well-informed choice and select an appropriate variant it is necessary to have a good understanding of a given method’s capabilities and limitations. Hence, developing this understanding is of central importance. Knowledge of the limitations and failure modes further has the advantage of suggesting ways to improve upon.

We will focus on one specific class of approximate inference methods. That is the class of message passing algorithms. As discussed in Section 2.5.1 message passing algorithms efficiently perform exact inference on tree-structured models; the existence of loops, however, has a severe effect on message passing and impedes the straightforward generalization (cf. Example 3). Nonetheless, the elegance and simplicity of performing only local operations remain indisputable

and it is tempting, therefore, to ignore the existence of loops, apply the same principles, and hope for a reasonable outcome.

Indeed, various approximate inference methods work according to this very principle. In fact, because of their appealing simplicity, similar concepts were independently introduced multiple times in different fields: Judea Pearl introduced message passing algorithms in the machine learning- and statistics-community [Pea88] for tree-structured graphs, terming it *belief propagation*. He already advocated the extension to loopy graphs as an approximate method [Pea88, Section 4.4] – nowadays often referred to as *loopy belief propagation* to emphasize the approximate nature of the method.¹² In the information theory community, message passing algorithms were first introduced in the PhD Thesis of Robert Gallager [Gal68] for decoding low-density-parity-check codes. The capabilities of message passing – known as the sum-product-algorithm – were, however, largely overlooked until the success of Turbo codes [BGT93]. Relying on the very same principles, this renewed interest finally put the original work into perspective and ultimately led to its well-deserved recognition. Very similar concepts are also applied with success in the signal-processing community as in the Kalman filter [Kal60] for state estimation, or in the Viterbi algorithm [Vit67] for hidden Markov models. Probably the first explicit application of local operations for approximating the global behavior has to be attributed to the physics community: Hans Bethe and Rudolf Peierls introduced this concept to the field of statistical physics with the aim of understanding the behavior of large, otherwise incomprehensible, interacting systems [Bet35, Pei36]. This approximation is known as the Bethe-Peierls approximation or the Cavity method.¹³

3.2 Preliminaries

Here in this section, we will finally define belief propagation. Let us consider a binary pairwise graphical model $\mathcal{U} = (\mathcal{G}, \Psi)$. First, we define messages that are passed along the edges, where the message from X_i to X_j is denoted by $\mu_{ij}^n(x_j)$ with $n \in \mathbb{Z}$ numbering the current iteration.¹⁴ The messages are updated according to the recursive rule:

$$\mu_{ij}^{n+1}(x_j) \propto \sum_{x_i \in \mathcal{X}} \Phi(x_i, x_j) \Phi(x_i) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(x_i). \quad (3.1)$$

To compute the messages, BP collects all messages sent to X_i , except from X_j and multiplies this product with the local potential $\Phi(x_i)$ and the pairwise potential $\Phi(x_i, x_j)$. Finally, the sum over all states $x_i \in \mathcal{X}$ is sent. In practice, the messages require some form of normalization [IFW05]; we will normalize the messages by $\alpha_{ij}^n \in \mathbb{R}_+^*$ so that $\sum_{x_j \in \mathcal{X}} \mu_{ij}^n(x_j) = 1$, which gives the messages a probabilistic interpretation. Note that all messages are consequently restricted to $\mu_{ij}(x_j) \in [0, 1]$ and the application of the update rule (3.1) does not change this.

Lemma 1. *Normalized messages, sent from node X_i to X_j over $(i, j) \in \mathbf{E}$, represent probabilities and remain so under successive application of the BP update equation – provided all messages are initialized to be positive.*

Proof. Positive potentials in (3.1) guarantee that all messages remain positive at every iteration.

¹² We will overload the terminology and, for the sake of brevity, always refer to it as belief propagation; irrespective whether the graph contains loops or not.

¹³ Nowadays, the connections between the coding- and the artificial intelligence community [KFL01] as well as the connections to the physics community [YFW05] are well established and may seem rather obvious in hindsight. Looking at the original literature, however, it becomes clear that, because of the different formalism, these findings were indeed rather surprising at first.

¹⁴ Two neighboring nodes are joined by a single edge by definition. Note, that information has to flow into both directions and that two distinct messages are passed along opposing directions over every edge.

Consequently, a normalization term α_{ij}^n exists so that $\sum_{x_j \in \mathcal{X}} \mu_{ij}^{n+1}(x_j) = 1$. \square

The update equation in (3.1) closely resembles the recursive definition of the messages in Section 2.5.1, except for one fundamental difference. For tree-structured graphs the recursive definition of the messages results from simply reorganizing the summations. Consequently, the messages converge once all messages were passed up and down the tree. For loopy graphs, however, such a reorganization is not possible – since the recursive definition of a message, say $\mu_{ij}(x_j)$, contains the message $\mu_{ji}(x_i)$ itself. For the very same reason, the notion of passing the messages up and down the tree is simply not possible anymore. Instead, one can neglect the existence of loops, update the message according to the recursive definition, accept the approximate nature of this approach, and hope for the messages to converge to the *fixed point messages* $\mu_{ij}^\circ(x_j)$.

After convergence of BP, one can approximate the marginals similar as for the tree-structured models. Note that the marginals were defined over any possible subset of random variables in (2.2), though we will only consider the singleton marginals $P_{X_i}(x_i)$ and the pairwise marginals $P_{X_i, X_j}(x_i, x_j)$ in this thesis. In general, the marginals over any subset of random variables are computed by the product of incoming messages times the associated potentials. Accordingly, we approximate the singleton and pairwise marginals by

$$\tilde{P}_{X_i}(x_i) = \frac{1}{Z_i} \Phi(x_i) \prod_{X_k \in \partial(i)} \mu_{ki}^\circ(x_i), \quad (3.2)$$

$$\tilde{P}_{X_i, X_j}(x_i, x_j) = \frac{1}{Z_{ij}} \Phi(x_i) \Phi(x_j) \Phi(x_i, x_j) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^\circ(x_i) \prod_{X_l \in \{\partial(j) \setminus X_i\}} \mu_{lj}^\circ(x_j), \quad (3.3)$$

where $Z_i, Z_{ij} \in \mathbb{R}_+^*$ guarantee that all probabilities sum to one. We denote the set of all approximated singleton- and pairwise marginals by

$$\tilde{P}_B = \{\tilde{P}_{X_i}(x_i), \tilde{P}_{X_i, X_j}(x_i, x_j) : X_i \in \mathbf{X}, (i, j) \in \mathbf{E}\}, \quad (3.4)$$

and refer to \tilde{P}_B as the pseudomarginals.¹⁵ This naming-convention should highlight the fact that, in general, BP only approximates the marginals. In fact, one may end up with pseudomarginals that do not correspond to any valid distribution at all.

Example 4 (Unrealizable Pseudomarginals).

Let us consider the pairwise graphical model in Figure 3.1 where \mathbf{X} is a set of three binary random variables with $x_i \in \mathcal{X} = \{-1, 1\}$.

The model is frustrated and contains two attractive edges with $J_{ij} = 1$ (depicted by solid lines) and one repulsive edge with $J_{ij} = -1$ (depicted by the dashed line).

Let us specify the pairwise potentials so that

$$\begin{aligned} \Phi_{X_1, X_2}(x_1, x_2) &= \begin{bmatrix} \exp(J) & \exp(-J) \\ \exp(-J) & \exp(J) \end{bmatrix}, \\ \Phi_{X_2, X_3}(x_2, x_3) &= \begin{bmatrix} \exp(J) & \exp(-J) \\ \exp(-J) & \exp(J) \end{bmatrix}, \\ \Phi_{X_3, X_1}(x_3, x_1) &= \begin{bmatrix} \exp(-J) & \exp(J) \\ \exp(J) & \exp(-J) \end{bmatrix}. \end{aligned}$$

¹⁵ Note that the pseudomarginals are also often called *beliefs*, hence the name belief propagation.




Figure 3.1: Example of a frustrated model. The solid edges are attractive and the dashed edge is repulsive; note that the product of couplings along all edges (cf. (2.18)) equates to a negative number, which characterizes the existence of frustrations. While X_2 tries to pull X_1 and X_3 towards the same state, both variables are pushed apart by each other because of the repulsive edge.

We further set all local potentials to identical values and choose $\theta = 0$ so that

$$\Phi_{X_i}(x_i) = \begin{bmatrix} \exp(\theta) \\ \exp(-\theta) \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \quad (3.5)$$

Belief propagation converges to a unique fixed point defined by $\mu_{ij}^\circ(x_j) = 0.5$ for all $x_j \in \mathcal{X}$. It is well-established that this particular fixed point exists and does not depend on the pairwise potentials if $\theta = 0$ (cf. Section 5.7.1). The fact that these messages $\mu_{ij}^\circ(x_j) = 0.5$ constitute a fixed point becomes immediately obvious if we plug in all potentials into the update equation (3.1).

The pairwise marginals are subsequently approximated by $\tilde{P}_{X_i, X_j}(x_i, x_j)$; in particular we have

$$\tilde{P}_{X_1, X_2}(x_1, x_2) = \begin{bmatrix} 0.44 & 0.06 \\ 0.06 & 0.44 \end{bmatrix}, \quad (3.6)$$

$$\tilde{P}_{X_2, X_3}(x_2, x_3) = \begin{bmatrix} 0.44 & 0.06 \\ 0.06 & 0.44 \end{bmatrix}, \quad (3.7)$$

$$\tilde{P}_{X_3, X_1}(x_3, x_1) = \begin{bmatrix} 0.06 & 0.44 \\ 0.44 & 0.06 \end{bmatrix}. \quad (3.8)$$

Now we will proof by contradiction that the pairwise marginals from above are unrealizable; i.e., one cannot specify a joint distribution over three random variables that has its pairwise marginals correspond to (3.6) - (3.8). To make this more precise, note that by (3.6) and by the sum-rule we have $\tilde{P}_{\mathbf{X}}(-1, +1, -1) + \tilde{P}_{\mathbf{X}}(-1, +1, +1) = 0.06$, so that the joint probability for both configurations must satisfy $\tilde{P}_{\mathbf{X}}(-1, +1, -1) \stackrel{!}{\leq} 0.06$ and $\tilde{P}_{\mathbf{X}}(-1, +1, +1) \stackrel{!}{\leq} 0.06$.

Accordingly, (3.6) - (3.8) reveal the joint probabilities for all configurations $\mathbf{x} \in |\mathcal{X}|^N$. In this example all joint probabilities must satisfy

$$\tilde{P}_{\mathbf{X}}(\mathbf{x}) \stackrel{!}{\leq} 0.06. \quad (3.9)$$

The total probability is thus bounded from above according to

$$\sum_{\mathbf{x}} \tilde{P}_{\mathbf{X}}(\mathbf{x}) \leq 2^N \cdot 0.06 = 0.48, \quad (3.10)$$

which violates the fundamental rules of probability. Consequently, the pseudomarginals obtained by BP cannot belong to a valid joint distribution.

The above example demonstrated why the pseudomarginals may fail to represent a valid distribution. Still, unrealizable pseudomarginals need not pose a serious problem. In fact, we usually resort to BP whenever evaluating the exact marginals is not possible, so that approximating the marginals reasonably well is often sufficient (even if they are unrealizable). The main motivation for considering BP was the hope for its efficient nature to carry over to loopy graphs. This is unfortunately not always the case and, in the presence of loops, BP – besides only approximating the marginals – may fail to converge altogether.

Nonetheless, despite all these shortcomings, BP often works surprisingly well, even for models that contain many loops. So far we have depicted BP as a pure heuristic without any theoretical motivation. This is unsatisfactory, and the situation was exactly like this for many years until a close relation to variational methods was revealed. The variational perspective of BP is appealing for several reasons: it relates to a similar problem, well studied in statistical physics; it provides a more lucid justification of BP; and it explains the approximate nature of BP. It is for all those reasons that the variational perspective spurred much of the research on BP in the last decade and significantly improved the theoretical understanding of BP.

3.3 Variational Interpretation of BP – A Physics Perspective

Ultimately, the aim of approximate inference is to approximate some unknown complex distribution in a way that simplifies answering probabilistic queries, as for example estimating the marginals. One particularly powerful way of approximating distributions is available in the form of variational methods. The variational free energy approach [JGJS99] was first formalized in the context of statistical mechanics by Richard Feynman [Fey72, Section 3.4]. Although variational methods originate from the physics literature, they are not limited to problems arising in this context and have become a reliable tool in the machine learning community. For our purpose, we focus on the Bethe free energy $\mathcal{F}_{\mathcal{B}}$ that bridges the gap from BP to variational methods.

We show how the Bethe free energy emerges quite naturally if we approximate $P_{\mathbf{X}}(\mathbf{x})$ according to the variational free energy principle and follow the excellent deductive presentation of [YFW05] or [Mac03, Chapter 33]. An exhaustive treatment of variational representations in the context of probabilistic graphical models that elucidates how a variety of approximate inference methods can be understood in terms of their variational representations is presented in the excellent review article [WJ08]. Despite its origin, we refrain, however, from discussing the insights of specific relevance for problems in physics and refer the interested reader to some of the many well-written books [Geo11, Hua63] for a more in-depth treatment of the Bethe approximation from the physics perspective.

When approximating a joint distribution $P_{\mathbf{X}}(\mathbf{x})$, we first introduce a *trial distribution* $\tilde{P}_{\mathbf{X}}(\mathbf{x})$. The aim is to manipulate the trial distribution so that the mismatch between both distributions is reduced. Therefore, we quantify the mismatch by the Kullback-Leibler divergence, defined according to

$$D(\tilde{P}_{\mathbf{X}}||P_{\mathbf{X}}) = \sum_{\mathbf{x} \in \mathcal{X}^N} \tilde{P}_{\mathbf{X}}(\mathbf{x}) \log \frac{\tilde{P}_{\mathbf{X}}(\mathbf{x})}{P_{\mathbf{X}}(\mathbf{x})}, \quad (3.11)$$

where, with slight abuse of notation, $\log(\cdot)$ corresponds to the natural logarithm. Note that the Kullback-Leibler is a non-symmetric function and not a proper distance measure. It does, however, satisfy the Gibbs inequality, i.e., it satisfies $D(\tilde{P}_{\mathbf{X}}||P_{\mathbf{X}}) \geq 0$ with equality if and only if $\tilde{P}_{\mathbf{X}}(\mathbf{x}) = P_{\mathbf{X}}(\mathbf{x})$.¹⁶

¹⁶ Different proofs can be found in the literature for this property of the Kullback-Leibler divergence that, e.g., rely on the log-sum or the Jensen's inequality [MM09, p.7].

Let $P_{\mathbf{X}}(\mathbf{x})$ belong to an exponential family, defined by (2.14), and let $\tilde{P}_{\mathbf{X}}(\mathbf{x})$ be an arbitrary trial distribution. Then we can rewrite the Kullback-Leibler divergence according to

$$D(\tilde{P}_{\mathbf{X}}||P_{\mathbf{X}}) = \sum_{\mathbf{x} \in \mathcal{X}^N} \tilde{P}_{\mathbf{X}}(\mathbf{x}) \log \tilde{P}_{\mathbf{X}}(\mathbf{x}) + \sum_{\mathbf{x} \in \mathcal{X}^N} \tilde{P}_{\mathbf{X}}(\mathbf{x}) E(\mathbf{x}) + \sum_{\mathbf{x} \in \mathcal{X}^N} \tilde{P}_{\mathbf{X}}(\mathbf{x}) \log \mathcal{Z}. \quad (3.12)$$

Overloading our notation, we further define the average energy $E(\tilde{P}_{\mathbf{X}}) = \mathbb{E}_{\mathbf{X}}(E(\mathbf{x}))$ and the entropy $S(\tilde{P}_{\mathbf{X}}) = \mathbb{E}_{\mathbf{X}}(-\log(\tilde{P}_{\mathbf{X}}(\mathbf{x})))$ so that

$$E(\tilde{P}_{\mathbf{X}}) = \sum_{\mathbf{x} \in \mathcal{X}^N} \tilde{P}_{\mathbf{X}}(\mathbf{x}) E(\mathbf{x}) \quad (3.13)$$

$$S(\tilde{P}_{\mathbf{X}}) = - \sum_{\mathbf{x} \in \mathcal{X}^N} \tilde{P}_{\mathbf{X}}(\mathbf{x}) \log \tilde{P}_{\mathbf{X}}(\mathbf{x}). \quad (3.14)$$

We can now plug the average energy and the entropy into the definition of the Kullback-Leibler divergence and express (3.12) in a simplified way according to

$$D(\tilde{P}_{\mathbf{X}}||P_{\mathbf{X}}) = \log \mathcal{Z} + E(\tilde{P}_{\mathbf{X}}) - S(\tilde{P}_{\mathbf{X}}). \quad (3.15)$$

Let us define two important quantities, namely the Helmholtz free energy (or the negative log-partition function)

$$\mathcal{F}_{\mathcal{H}} = -\log \mathcal{Z}, \quad (3.16)$$

and the Gibbs free energy, which is given by

$$\mathcal{F}_{\mathcal{G}}(\tilde{P}_{\mathbf{X}}) = E(\tilde{P}_{\mathbf{X}}) - S(\tilde{P}_{\mathbf{X}}). \quad (3.17)$$

Then, we can finally rearrange terms, plug the Gibbs free energy into (3.15) and express it in a sensible way that suggests how to estimate $\mathcal{F}_{\mathcal{H}}$ (and thus the partition function \mathcal{Z}).

$$\mathcal{F}_{\mathcal{G}}(\tilde{P}_{\mathbf{X}}) = \mathcal{F}_{\mathcal{H}} + D(\tilde{P}_{\mathbf{X}}||P_{\mathbf{X}}). \quad (3.18)$$

The properties of the Kullback-Leibler divergence imply that $\mathcal{F}_{\mathcal{G}} \geq \mathcal{F}_{\mathcal{H}}$ with equality if and only if $\tilde{P}_{\mathbf{X}}(\mathbf{x}) = P_{\mathbf{X}}(\mathbf{x})$. This is a pleasant formalism that opens the door for variational approaches; in particular minimizing $\mathcal{F}_{\mathcal{G}}(\tilde{P}_{\mathbf{X}})$ provides $\mathcal{F}_{\mathcal{H}}$ and additionally, at its minimum, where the Kullback-Leibler divergence vanishes, the trial distribution converges to the exact one $P_{\mathbf{X}}(\mathbf{x})$. More formally, for the set of all valid trial distribution $\tilde{P}_{\mathbf{X}}(\mathbf{x})$ over \mathcal{X}^N we have

$$P_{\mathbf{X}}(\mathbf{x}) = \arg \min_{\tilde{P}_{\mathbf{X}}(\mathbf{x})} \mathcal{F}_{\mathcal{G}}(\tilde{P}_{\mathbf{X}}). \quad (3.19)$$

Note that this constitutes the central problems of inference, i.e., evaluating the partition function and computing marginal distributions (cf. Section 2.5). The formulation as an optimization problem provides an elegant way of solving those inference problems but – in its current form – is of limited relevance. Evaluation of the Gibbs free energy alone becomes infeasible with increasing model size as both terms in the Gibbs free energy require a summation over exponentially many terms (cf.(3.13) and (3.14)). The importance of (3.19), however, lies in the fact that a whole family of tractable approximation methods results from restricting the possible choices of trial distributions.

One popular approximation method that follows this idea is the *mean field* method. In its simplest form, one assumes that the trial distribution is defined over independent random variables. Consequently, the joint distribution factorizes into the product of the singleton marginals

according to

$$\tilde{P}_{\mathbf{X}}(\mathbf{x}) = \prod_{X_i \in \mathbf{X}} \tilde{P}_{X_i}(x_i). \quad (3.20)$$

Restricting ourselves to trial distributions of the form (3.20), the minimization of the Gibbs free energy suddenly becomes tractable. The minimizer of (3.19) then approximates the marginals. Some problems exist for which the mean field approximation even becomes asymptotically exact, as for example for infinite size Ising grid graphs with $\hat{d}_{\mathcal{G}} = \infty$ [MM09, p.80]. The approximation tends to require a high average degree to work well. In general, however, this is not the case and the mean field method often provides poor approximations. This is mainly because of statistical dependencies that are present in $P_{\mathbf{X}}(\mathbf{x})$, which the mean field method fails to account for.

It is an obvious next step to enhance the approximation quality by accounting for the correlations between pairs of random variables. One prevalent way comes in the form of the Bethe approximation that is of particular relevance for this work. Not only does the Bethe free energy $\mathcal{F}_{\mathcal{B}}$ restrict the class of trial distributions but it inherently approximates the Gibbs free energy as well. The latter is a result of the average energy $E_{\mathcal{B}}(\tilde{P}_{\mathcal{B}})$ and the Bethe entropy $S_{\mathcal{B}}(\tilde{P}_{\mathcal{B}})$ that are evaluated over the pseudomarginals $\tilde{P}_{\mathcal{B}}$ instead of the joint distribution according to

$$\begin{aligned} E_{\mathcal{B}}(\tilde{P}_{\mathcal{B}}) &= - \sum_{\mathbf{x}_m \in \{X_i \in \mathbf{X}\} \cup \{X_i, X_j : (i,j) \in \mathbf{E}\}} \tilde{P}_{\mathbf{x}_m}(\mathbf{x}_m) \cdot \ln \Phi_{\mathbf{x}_m}(\mathbf{x}_m) \\ &= - \sum_{X_i \in \mathbf{X}} \sum_{x_i \in \mathcal{X}} \tilde{P}_{X_i}(x_i) \ln \Phi(x_i) - \sum_{(i,j) \in \mathbf{E}} \sum_{x_i, x_j \in \mathcal{X}^2} \tilde{P}_{X_i, X_j}(x_i, x_j) \ln \Phi(x_i, x_j) \end{aligned} \quad (3.21)$$

$$\begin{aligned} S_{\mathcal{B}}(\tilde{P}_{\mathcal{B}}) &= - \sum_{(i,j) \in \mathbf{E}} \sum_{x_i, x_j \in \mathcal{X}^2} \tilde{P}_{X_i, X_j}(x_i, x_j) \ln \tilde{P}_{X_i, X_j}(x_i, x_j) \\ &\quad + \sum_{X_i \in \mathbf{X}} (d_i - 1) \sum_{x_i \in \mathcal{X}^2} \tilde{P}_{X_i}(x_i) \ln \tilde{P}_{X_i}(x_i). \end{aligned} \quad (3.22)$$

This subsequently defines the Bethe free energy in accordance with (3.17) so that

$$\mathcal{F}_{\mathcal{B}}(\tilde{P}_{\mathcal{B}}) = E_{\mathcal{B}}(\tilde{P}_{\mathcal{B}}) - S_{\mathcal{B}}(\tilde{P}_{\mathcal{B}}) \quad (3.23)$$

$$\begin{aligned} &= \sum_{(i,j) \in \mathbf{E}} \sum_{x_i, x_j} \tilde{P}_{X_i, X_j}(x_i, x_j) \ln \frac{\tilde{P}_{X_i, X_j}(x_i, x_j)}{\Phi(x_i, x_j)} - \sum_{X_i} \sum_{x_i} \tilde{P}_{X_i}(x_i) \ln \Phi(x_i) \\ &\quad - \sum_{X_i} (d_i - 1) \sum_{x_i} \tilde{P}_{X_i}(x_i) \ln \tilde{P}_{X_i}(x_i). \end{aligned} \quad (3.24)$$

In comparison to the Gibbs free energy, the Bethe free energy reduces the computational burden drastically; solely because $\mathcal{F}_{\mathcal{B}}$ is evaluated over the pseudomarginals, i.e., the singleton and pairwise marginals, instead of the full joint distribution. Just as the minimum of $\mathcal{F}_{\mathcal{G}}$ provides the partition function and the marginals, one would hope that – as $\mathcal{F}_{\mathcal{B}}$ approximate $\mathcal{F}_{\mathcal{G}}$ – minimizing $\mathcal{F}_{\mathcal{B}}$ will provide approximations to the partition function and the marginals.

In order to minimize $\mathcal{F}_{\mathcal{G}}$ efficiently, we had to restrict the choice of trial distributions. A similar restriction is required for minimizing $\mathcal{F}_{\mathcal{B}}$; the most obvious class of trial distributions are realizable pseudomarginals that adhere to the sum-rule of probability. This constraints the pseudomarginals so that they correspond to some valid distribution $P_{\mathbf{X}}(\mathbf{x})$; we refer to this set as the marginal polytope.

Definition 3.3.1 (Marginal Polytope). The marginal polytope is the set of all pseudomarginals that are jointly realizable by a valid joint distribution $P_{\mathbf{X}}(\mathbf{x})$; i.e.,

$$\mathbb{M}(\mathcal{G}) = \left\{ \tilde{P}_{X_i}(x_i), \tilde{P}_{X_i, X_j}(x_i, x_j) : \begin{aligned} \tilde{P}_{X_i}(x_i) &= \sum_{\mathbf{x} \setminus X_i} P_{\mathbf{X}}(\mathbf{x}) : X_i \in \mathbf{X}, \\ \tilde{P}_{X_i, X_j}(x_i, x_j) &= \sum_{\mathbf{x} \setminus \{X_i, X_j\}} P_{\mathbf{X}}(\mathbf{x}) : (i, j) \in \mathbf{E} \end{aligned} \right\}. \quad (3.25)$$

Note that the number of constraints in the marginal polytope depends on the structure of the graph as the singleton marginals and the pairwise marginals are only defined over the set of all nodes and the set of all edges respectively. We will refrain from making the dependence on the graph explicit, however, and only refer to the marginal polytope as \mathbb{M} .

The constrained minimization of the Bethe free energy, i.e., $\min_{\mathbb{M}} \mathcal{F}_{\mathcal{B}}(\tilde{P}_B)$, seems like a standard optimization problem. It is, but the huge amount of constraints in \mathbb{M} impedes the optimization in practice [Mur12, Section 22.3]. This issue is usually dealt with by relaxing the set of constraints. Accordingly we may relax the requirement from *globally realizable marginals* to *locally consistent marginals*. This makes the dependence on $P_{\mathbf{X}}(\mathbf{x})$ superfluous and only requires two properties to be satisfied: all singleton and pairwise marginals have to be properly normalized, and the singleton marginals must be consistent with the pairwise marginals. These constraints then define the local polytope \mathbb{L} .

Definition 3.3.2 (Local Polytope). The local polytope is the set of all pseudomarginals that are locally consistent, i.e.,

$$\mathbb{L}(\mathcal{G}) = \left\{ \tilde{P}_{X_i}(x_i), \tilde{P}_{X_i, X_j}(x_i, x_j) : \begin{aligned} \sum_{x_i \in \mathcal{X}} \tilde{P}_{X_i}(x_i) &= 1, \\ \tilde{P}_{X_i}(x_i) &= \sum_{x_j \in \mathcal{X}} \tilde{P}_{X_i, X_j}(x_i, x_j) \end{aligned} \right\} \quad (3.26)$$

As for the marginal polytope, we will also refrain from making the dependence on the graph structure explicit and only refer to the local polytope by \mathbb{L} . Further note that \mathbb{L} contains fewer constraints than \mathbb{M} and thus provides an outer bound on the marginal polytope (with equality for tree-structured models (cf. [WJ08, Prop. 4.1])).

It is worth pointing out a few differences that set the Bethe approximation apart from other variational methods. On the one hand, we would like to stress that the Bethe approximation approximates the marginals well in many cases. On the other hand, there are some conceptual shortcomings: First, the minimization must be performed over the local polytope; the obtained marginals are consequently only locally consistent and, in general, do not belong to a valid joint distribution. Second, except for certain sub-classes of graphical models, the Bethe free energy does not upper bound (nor lower bound) the original objective. This seems to fundamentally jeopardize the main assumption that the minimum of $\mathcal{F}_{\mathcal{B}}(\tilde{P}_B)$ should be somewhat close to the minimum of $\mathcal{F}_{\mathcal{G}}(\tilde{P}_{\mathbf{X}})$.

To summarize, the Bethe approximation has the potential to work well and may thus serve as an efficient approximation method but, at the same time, its properties indicate that performance guarantees will often be hard to come by. This highlights the need for developing a good understanding of the Bethe approximation and its shortcomings. The Bethe approximation is well-studied in statistical physics. Note, however, that the models studied in physics differ from the ones studied in computer science in one important aspect. Whereas physicists usually study the behavior in the thermodynamic limit, i.e., for graphs of infinite size,¹⁷ in computer science,

¹⁷ This allows one to take limits and admits an analytical treatment in some cases.

we are confronted with models of finite size. Although we expect a similar behavior to a certain degree, we will see that finite-size effects will often have a notable effect on the properties of \mathcal{F}_B .

3.3.1 Energy Landscape of the Bethe Free Energy

Let us now focus on the practical aspects of minimizing the Bethe free energy directly. In particular, we will take a closer look at how to approximate the marginals and the partition function in this setting. Similar as the partition function corresponds to the Helmholtz free energy (cf. (3.16)), we introduce the *Bethe partition function* \mathcal{Z}_B that corresponds to the Bethe free energy according to

$$\mathcal{Z}_B = \exp(-\mathcal{F}_B). \quad (3.27)$$

Just as \mathcal{F}_B approximates \mathcal{F}_H , \mathcal{Z}_B approximates \mathcal{Z} . Note that there is literally no difference between both quantities and instead of minimizing \mathcal{F}_B one could equivalently maximize \mathcal{Z}_B . We stick to the notion of minimizing the Bethe free energy, primarily for historical reasons.

Let us recap the two key-ingredients of the Bethe approximation again: we must replace the Gibbs free energy by the Bethe free energy and replace the constraints of the marginal polytope by the ones of the local polytope. This gives us the global minimum of the constrained Bethe free energy according to

$$\mathcal{F}_B^* = \min_{\mathbb{L}} \mathcal{F}_B(\tilde{P}_B), \quad (3.28)$$

where the pseudomarginals are specified by the associated minimizer

$$\tilde{P}_B^* = \arg \min_{\mathbb{L}} \mathcal{F}_B(\tilde{P}_B). \quad (3.29)$$

Although the relaxation to \mathbb{L} reduces the complexity, it concurrently alters the energy landscape with drastic consequences. As opposed to the convex Gibbs free energy, the Bethe free energy is generally a non-convex function (cf. Example 5).

We will denote all stationary points of the constrained Bethe free energy by

$$\mathcal{F}_B^\circ \in \{\mathcal{F}_B : \nabla \mathcal{F}_B(\tilde{P}_B) = 0\}. \quad (3.30)$$

Finally, local minima of the constrained Bethe free energy are particularly relevant. Let us express the Hessian of the Bethe free energy by $\nabla^2 \mathcal{F}_B(\tilde{P}_B)$. Then we denote local minima explicitly by

$$\mathcal{F}_B^m \in \{\mathcal{F}_B : \nabla \mathcal{F}_B(\tilde{P}_B) = 0, \nabla^2 \mathcal{F}_B(\tilde{P}_B) \text{ is positive definite}\}. \quad (3.31)$$

Note, that the different types of stationary points consequently satisfy $\mathcal{F}_B^* \subseteq \{\mathcal{F}_B^m\} \subseteq \{\mathcal{F}_B^\circ\}$.

3.3.2 The Bethe Approximation and Belief Propagation

One might ask where the added value of the variational approach lies exactly? The main reason is the immediate connection between the stationary points of the Bethe free energy and the fixed points of BP. The rich history of studying the Bethe free energy thus benefits our understanding of BP directly, as the theoretical properties carry over through this relationship.

Explicit Connection

There is a fundamental connection between the stationary points of the Bethe free energy and the BP fixed points μ° (and the associated pseudomarginals \tilde{P}_B°). In a nutshell, the stationary points are in a one-to-one correspondence with the fixed points. We take advantage of the fact that \mathcal{F}_B is defined over the pseudomarginals (cf. (3.28)) and express the relation between the stationary points \mathcal{F}_B° and the pseudomarginals according to

$$\mathcal{F}_B^\circ = \mathcal{F}_B(\tilde{P}_B^\circ). \quad (3.32)$$

This correspondence becomes apparent when specifying the minimization of \mathcal{F}_B explicitly and including the constraints of the local polytope as a Lagrangian. Taking the derivatives and setting them to zero then yields the BP update equations [YFW01]. Consequently, at stationary points of \mathcal{F}_B (where all partial derivatives are zero), all BP messages remain unaffected by the update rule (which constitutes a BP fixed point). In fact, every fixed point of BP is an interior stationary point of the constrained Bethe free energy [YFW05]. Note that the nature of this correspondence further reveals a subtle, yet important, detail of BP: updating a message forces one gradient at a time to zero while keeping all other variables fixed, which albeit often going downwards is not a gradient descent step per se. The variational principle, however, suggests that we should specifically consider minima of \mathcal{F}_B ; luckily, the BP updates still tend to proceed in a sensible way towards minima of \mathcal{F}_B [Hes03, AM00].

Implications

Stable fixed points of BP (see. Section. 4.4 for a thorough discussion on stability) are of particular relevance in practice. Let us index all stable fixed points $\tilde{P}_B^{(s)}$ by $s = 1, \dots, S$. Every stable fixed point (s) then has an associated local minimum $\mathcal{F}_B^{(s)}$, an associated partition function $\mathcal{Z}_B^{(s)}$, and associated pseudomarginals $\tilde{P}_B^{(s)}$; we denote the set of all S stable fixed points by

$$\mathbf{S} = \left\{ (\mathcal{Z}_B^{(1)}, \tilde{P}_B^{(1)}), \dots, (\mathcal{Z}_B^{(S)}, \tilde{P}_B^{(S)}) \right\}. \quad (3.33)$$

Likewise, we consider the set of all fixed points that constitute minima of the Bethe free energy

$$\mathbf{M} = \left\{ (\mathcal{Z}_B^1, \tilde{P}_B^1), \dots, (\mathcal{Z}_B^M, \tilde{P}_B^M) \right\}, \quad (3.34)$$

and the total set of fixed points

$$\mathbf{T} = \left\{ (\mathcal{Z}_B^1, \tilde{P}_B^1), \dots, (\mathcal{Z}_B^T, \tilde{P}_B^T) \right\}. \quad (3.35)$$

Further note that all stable fixed points of BP must be minima of \mathcal{F}_B . In the presence of frustrated cycles, however, minima may be unstable as well [Hes03, MWJ99]. The different types of fixed points thus relate to each other according to $\mathbf{S} \subseteq \mathbf{M} \subseteq \mathbf{T}$.

One important question is to understand under which conditions a fixed point is unique, i.e., when $|\mathbf{T}| = 1$. It seems appealing to utilize the connection between BP and \mathcal{F}_B for that purpose. Besides for very simple models, however, conditions for convexity of \mathcal{F}_B are hard to come by, and established conditions are often far from necessary [PA02, Hes04]. A good overview of different conditions for convexity of \mathcal{F}_B is presented in [MK07].

Besides these results for uniqueness, not many results aim to characterize the expected number of fixed points. One work that pursues this direction is [WF09]. Note that – at least as long as some proper form of message normalization is used [MLF11] – the number of fixed points is always finite (cf. Theorem 7 and Lemma 2 in [WF09]). Another important insight is that the number of fixed points is always odd.

We will later characterize the number of fixed points and assess their stability in Chapter 5 and in Chapter 7 for a range of models. As of now, we will show the correspondence between stationary points of \mathcal{F}_B and fixed points of BP for one exemplary model.

Example 5 (Energy Landscape of \mathcal{F}_B for an Attractive Ising Model).

Let us take a closer look at the relationship between the Bethe free energy and BP by means of an attractive Ising model, specified on an infinite-size two-dimensional grid graph. Let all variables have identical local potentials, specified by $\theta_i = \theta > 0$, and all edges have identical pairwise potentials, specified by $J_{ij} = J > 0$. We illustrate a slice of the Bethe- and the Gibbs- free energy along the marginals of one variable $P_{X_i}(x_i)$ for a model with weak couplings (Figure 3.2) and for a model with strong couplings (Figure 3.3) Note that the Bethe free energy upper bounds the Gibbs free energy for attractive models (cf. Section 3.4).




Figure 3.2: Schematic illustration of the energy landscape that shows a slice along $P_{X_i}(x_i)$ for an attractive model with weak couplings. The exact and the approximate results are depicted for both the partition function and the marginals.

If the couplings are sufficiently small, \mathcal{F}_B is convex and BP has a unique and stable fixed point that has the whole message-space as a region of attraction (cf. Section 5.6). BP will thus always converge independently of its initialization.




Figure 3.3: Schematic illustration of the energy landscape for an attractive model with strong couplings. The Bethe free energy is non-convex and BP may converge to one of the local minima. Note how $\tilde{P}_{X_i}^2(x_i)$ is more accurate than $\tilde{P}_{X_i}^1(x_i)$.

Let us now consider a model with strong couplings. Note how $\mathcal{F}_{\mathcal{B}}$ becomes non-convex and exhibits multiple stationary points that correspond to BP fixed points. We will later show in Chapter 5 that both local minima are stable fixed points and that the message initialization will determine to which fixed point BP converges. This has a major influence on the overall quality of the approximated quantities as the difference to the exact marginals $P_{X_i}(x_i)$ (i.e., the values minimizing the Gibbs free energy) and thus the accuracy varies considerably between different fixed points.

3.3.3 Other Variational Approaches

The one-to-one correspondence between stationary points of $\mathcal{F}_{\mathcal{B}}$ and fixed points of BP led to an improved understanding of BP and paved the way for methods that minimize the Bethe free energy directly [WT01, WT03, YR03]. The minimization of $\mathcal{F}_{\mathcal{B}}$, however, remains non-trivial and usually requires additional considerations to render it viable.

The consideration of the Bethe free energy opens the door for provable convergent algorithms. Belief optimization [WT01] minimizes $\mathcal{F}_{\mathcal{B}}$ by following the negative gradient and is guaranteed to converge. The minimization takes place along the edges of the local polytope, which guarantees that the marginalization constraints in (3.26) are satisfied throughout. Alternatively, one can decompose the non-convex $\mathcal{F}_{\mathcal{B}}$ into a convex and a concave problem (this decomposition is in general not unique) and optimize both objectives in an alternating fashion. This procedure, known as the constrained convex-concave procedure [YR03], is guaranteed to obtain a stationary point of $\mathcal{F}_{\mathcal{B}}$. Although both methods converge to stationary points of $\mathcal{F}_{\mathcal{B}}$, two major limitations remain: first, the approximation quality may vary considerably between different stationary points and only obtaining some (local) extremum may yield sub-optimal solutions; and second, despite the appeal of convergence guarantees, run-time guarantees are often just as important in practice.

To counteract some of these problems, one can also relax the Bethe free energy and come up with well-behaved surrogates that can be minimized efficiently. Convex surrogates seem to be specifically well-suited for this task and have thus received considerable attention. One can, for example, upper bound the objective by a convex function as in tree-reweighted belief propagation (TRW) [WJW03b, WJW05, Kol06] that enforces a concave entropy-term as a combination of tree-entropies. Note, however, that there are multiple ways of coming up with convex surrogates [MGW09, GJ07, HS08]. One comprehensive overview that unifies many different approaches in terms of the chosen counting number is presented in [MJGF09]. Overall these convex versions are well-behaved and can be optimized efficiently. A trade-off between convergence-properties and accuracy, however, persists and – if it can be minimized – the Bethe approximation often outperforms its convex surrogates in terms of accuracy [MJGF09, WJ14a].

This observation led to a renewed focus on trying to minimize the Bethe free energy efficiently (i.e., in polynomial run-time). The efficient minimization of $\mathcal{F}_{\mathcal{B}}$ becomes possible if we impose certain properties on the graphical model (in terms of its structure or parameters) and consider ϵ -approximating the stationary points. In particular, this includes sparse models [Shi12], where a projection scheme in the minimization task allows for a fully polynomial-time approximation on graphs with $\max(d_i) = \mathcal{O}(\log N)$. For attractive models (not necessarily sparse), this algorithm is further improved in [WJ14a] so that it obtains the global minimum of the Bethe free energy. If both properties are fulfilled, i.e., for locally tree-like attractive models the Bethe approximation is exact and can be optimized efficiently [DM10].

From a completely different point of view, the approximation quality can also be enhanced by better approximations of $\mathcal{F}_{\mathcal{G}}$. The concept of the Bethe approximation, that only accounts for the singleton- and pairwise marginals, generalizes to the Kikuchi method that accounts for the

marginals of larger cliques as well. The same principle – i.e., considering larger cliques – generalizes BP as well; this method is accordingly termed as generalized belief propagation [YFW01]. Both the accuracy and the convergence properties improve by adopting the computations to larger cliques. Although larger cliques inevitably increase the computational complexity, the principles of generalized belief propagation are flexible enough to allow moving freely along this trade-off.

Gauge transformations [CC06, CCT08] are somewhat similar in that they work with the exact partition function \mathcal{Z} directly while simplifying its estimation. This is achieved by expressing the exact partition function via a loop-series expansion of \mathcal{Z}_B . Interestingly, all terms in the loop-series correspond to fixed points of BP, which suggests one way of computing them. The number of terms in the expansion, however, may be large for models with many loops so that computing all terms is often not an option. Nonetheless, gauge transformations provide a principled approach to improve upon BP by at least accounting for some terms in the loop-series.

To conclude this section, a variety of variational approaches are available that aim to find the sweet spot between accuracy and complexity when performing approximate inference. Most of the above methods build upon the fundamental principles of the Bethe approximation, that – although providing a considerable simplification over the Gibbs free energy – remains problematic to minimize in practice.

3.4 Approximation Quality

So far we have already discussed various approximate inference methods. Different methods will perform differently and, depending on the given model, achieve varying accuracy. If we want to evaluate and compare approximate inference methods, it is important to measure the accuracy of the approximation. Having said that, approximate inference methods are usually applied to problems that forbid the computation of the exact solution. This prohibits measuring the accuracy by comparison to the exact solution and highlights the need for some qualitative measures of the expected performance for a given problem. Essentially, one would like to provide model-specific performance guarantees and bounds on the approximation error. Few bounds on the approximation quality are, however, established and some bounds require considerable computational resources on their own, not to mention that most bounds are often relatively loose.

In this section, we outline how the error in the approximation of the partition function and the marginals will be measured throughout the thesis. Moreover, we will discuss some of the available error-bounds for both quantities.

Partition Function

The error of the partition function is usually evaluated in terms of the relative error between the log-partition functions (i.e., the negative value of the Bethe free energy) according to

$$E_{\mathcal{Z}}(m) = \frac{|\log \mathcal{Z}_B^m - \log \mathcal{Z}|}{\log \mathcal{Z}} = \frac{|\mathcal{F}_{\mathcal{G}} - \mathcal{F}_B^m|}{-\mathcal{F}_{\mathcal{G}}}, \quad (3.36)$$

where $\mathcal{Z}_B^m = \mathcal{Z}_B(\tilde{P}_B^m)$ is the Bethe partition function of the m^{th} fixed point [GMK07]. Note that (3.36) quantifies the error in the partition function and in the free energy.

Existing bounds on the partition function usually combine an upper bound [WJW05, JJ97] with some lower bound as e.g., the naive mean field [WJ08]. Other bounds are based on the loop-series expansions [WSW08] or the non-backtracking operator [SKZ14].

One intriguing detail of the Bethe approximation is that it only approximates the partition function but neither provides an upper nor a lower bound of it. For the important class of attractive models, however, the Bethe partition function does indeed lower bound the partition function, i.e., $\mathcal{Z}_B < \mathcal{Z}$ [Ruo12]. This has the important consequence that minimizing \mathcal{F}_B is always optimal with respect to the approximation quality of the partition function for attractive models, since $\arg \min_{\mathcal{Z}_B^m} (E_{\mathcal{Z}}(m)) = \exp(-\min_{\mathbb{L}} \mathcal{F}_B(\tilde{P}_B^m))$.

An alternative, arguably more intuitive, proof that reveals how the Bethe partition function lower bounds the true partition function relies on the concept of clamping [WJ14b]. The main idea is to condition (i.e., to clamp) on a variable taking one specific value at a time and to evaluate the partition function as a sum over all sub-partition functions. This brings the advantage of working directly on the Bethe free energy and obviates the need for relying on additional concepts such as graph covers [Von13,Ruo13] or loop-series expansions [WSW08]. Moreover, for attractive models clamping always improves the approximation quality of the partition function over BP. This is particularly true if selecting the clamped variables wisely, as for example according to some heuristics [Wel16]. Clamping also improves the upper and lower bounds of \mathcal{Z} as it can be utilized to improve the partition function estimates of TRW and the mean field method [WD16].

Closely related to BP, but non-iterative, is the mini-bucket elimination scheme [DR03]. Mini-bucket elimination also provides upper and lower bounds on the partition function and offers a trade-off between the accuracy of the bound and computational efficiency based on the clique-size considered. One can further generalize this concept based on Hölders inequality [LI11] and efficiently compute bounds with good quality, where the improvement on the estimated bounds is because of incorporating the concept of TRW. The method of gauge transformations further generalized this concept and allows one to obtain tighter bounds on the partition function with similar computational effort [ACSW18].

Marginals

We measure the error of the singleton marginals by the mean squared error (MSE), where the error of the approximated marginals at the m^{th} fixed point is given by

$$E_P(m) = \frac{1}{N} \sum_{X_i \in \mathbf{X}} \|P_{X_i}(x_i) - \tilde{P}_{X_i}^m(x_i)\|_2^2 \quad (3.37)$$

$$= \frac{2}{N} \sum_{X_i \in \mathbf{X}} (P_{X_i}(x_i = 1) - \tilde{P}_{X_i}^m(x_i = 1))^2. \quad (3.38)$$

The simplification in (3.38) is because of symmetry properties for binary random variables. Note that one can replace the squared l_2 -norm in (3.37) by any other norm if desired; in particular the l_∞ -norm is considered by some authors to measure the worst-case error.

The expected mean describes the response of the system to the field θ [MM09] according to

$$\langle m \rangle = \mathbb{E}(m) = \frac{1}{N} \sum_{X_i \in \mathbf{X}} m_i = \frac{1}{N} \sum_{X_i \in \mathbf{X}} \mathbb{E}(X_i), \quad (3.39)$$

where we parameterize binary random variables by their mean (cf. (2.21)). Note that the difference between the expected mean of the exact marginals $\langle m \rangle$ and of the approximated marginals $\langle \tilde{m} \rangle$ is identical to the sum over all marginal errors

$$\langle m \rangle - \langle \tilde{m} \rangle = \frac{2}{N} \sum_{X_i \in \mathbf{X}} P_{X_i}(x_i = 1) - \tilde{P}_{X_i}(x_i = 1). \quad (3.40)$$

Note that we consider binary random variables with $\mathcal{X} = \{-1, 1\}$; if we would consider $\mathcal{X}' = \{0, 1\}$ instead, the expected mean would change according to $\langle m' \rangle = \frac{1}{N} \sum_{i=1}^N P_{X_i}(x_i = 1) = \frac{1}{2}(\langle m \rangle + 1)$.

Some methods quantify and bound the approximation error of the marginals instead of $E_{\mathcal{Z}}(m)$ by computing a confidence interval (i.e., an upper- and lower bound) on the exact marginals.¹⁸ Bound propagation [LK03] computes bounds for small clusters of nodes; similar to BP, these bounds are then propagated throughout the graph and provide bounds on the singleton marginals after convergence. The computational complexity is determined by the tree-width (similar as for the Junction tree algorithm), thus limiting the applicability to models with few loops.

Similar in principle, but more efficient in terms of computational complexity, is the recursive propagation of the bounds over sub-trees of the model [MK09]. The computational complexity is determined by the support of the random variables. This renders the method applicable to models with high connectivity as well.

From a completely different perspective, TRP provides bounds on the approximation error as well [WJW03a]. TRP considers spanning trees to approximate the marginals, which suggests computing bounds on the marginals over spanning trees as well. Regarding the computational complexity, there is one particularly intricate ingredient involved in bounding the error over the spanning trees, with the problem being the requirement for the log-partition function that cannot be computed exactly for complex models.

Some of these issues are resolved by computing the bounds over self-avoiding walk trees instead [Ihl07]. If we compare the quality of the bounds we observe that the proposed methods in [MK09] and [WJW03a] are of similar quality. The computation over self-avoiding walk trees impacts the quality of the obtained bounds: for models with weak interactions (where BP converges fast) the tightest bounds are provided by [Ihl07]; for models with strong interactions (where BP performs worse) the confidence intervals become wider than the ones provided by [WJW03a].

All those methods aim to estimate the expected performance of BP in terms of the marginal error and while this often estimates the performance well, all presented methods have one particular problem in common. Increasing the coupling strength degrades the quality of the estimated bounds.

3.5 BP as a Dynamical System

It is often convenient to consider BP as a dynamical system, i.e., as a discrete time map (see Definition 4.1.1 for a formal definition). Let $\boldsymbol{\mu}^n = \{\mu_{ij}^n(x_j), \mu_{ji}^n(x_i) : (i, j) \in \mathbf{E}\}$ be the set of all messages at iteration n , then we denote the mapping induced by the update equations of BP in (3.1) according to

$$\boldsymbol{\mu}^{n+1} = \mathcal{BP}(\boldsymbol{\mu}^n). \quad (3.41)$$

In accordance with the notion of fixed point messages $\mu_{ij}^\circ(x_j)$ that remain unaffected under the application of BP, we refer to the set of all fixed point messages by $\boldsymbol{\mu}^\circ$. We will often refer to $\boldsymbol{\mu}^\circ$ simply as fixed point and write

$$\boldsymbol{\mu}^\circ = \mathcal{BP}^\circ(\boldsymbol{\mu}), \quad (3.42)$$

where \mathcal{BP}° updates the messages until convergence.

¹⁸ Note that upper and lower bounds of the partition function can also be related to upper and lower bounds of the marginals [WJ14a].

Remember that BP is neither guaranteed to provide accurate marginals nor guaranteed to converge. Failure of BP (to provide accurate results) can be attributed to the existence of multiple (stable) fixed points or the existence of unstable fixed points (for which the messages oscillate far away from any fixed point [Wei00, MK07, IFW05]). Thus, if we want to make BP more robust, we should modify BP such that it converges towards an accurate stable fixed point. Potential modifications include variants that minimize \mathcal{F}_B directly or minimize alternative objectives as e.g., convex surrogates (cf. Section 3.3.3). From a completely different perspective, one can modify the update procedure directly – ignoring the effect on the variational function – so that the overall performance improves. In particular, the consideration of BP as a dynamical system comes with a flexible formalism that suggests multiple ways of modifying and enhancing BP.

3.5.1 Enhancing Belief Propagation’s Properties

BP, as presented in Section 3.1, is only one specific instance in a broad class of algorithms that fit under the umbrella of more general message passing methods. We will see that, although similar in nature, the choice of the respective message passing algorithm is not a mere formality but influences the behavior significantly – and is therefore of great practical relevance. The family of message passing algorithms is defined by a set of key-ingredients, where all message passing algorithms proceed along with the following three steps:

- **Initialization:** Each message $\mu_{ij}(x_j) \in \boldsymbol{\mu}$ is assigned an initial value $\mu_{ij}^1(x_j)$.
- **Update:** An update function is defined that acts on the set of messages $\mathcal{F} : \boldsymbol{\mu}^n \rightarrow \boldsymbol{\mu}^{n+1}$. This update function usually consists of one function per message $f_{ij} : \boldsymbol{\mu}^n \rightarrow \mu_{ij}^{n+1}(x_j)$.
- **Evaluation:** An evaluation or decision function maps the set of messages to the values of interest. In particular, we are interested in approximating the marginals (by the pseudomarginals) and in approximating the partition function (by the Bethe partition function); thus we consider evaluation functions for the pseudomarginals $\mathcal{D}_P : \boldsymbol{\mu} \rightarrow \tilde{P}_B$ and for the Bethe partition function $\mathcal{D}_Z : \boldsymbol{\mu} \rightarrow Z_B$.

This rather general formulation of message passing algorithms grants a certain degree of freedom in specifying an algorithm. Together, those aspects define the overall behavior of BP and every aspect may be tuned individually to enhance its performance. Not a single best algorithm exists though and, depending on the application, one or the other variant may perform better.

We will now discuss some established ways of enhancing BP’s performance. We particularly focus on how all those variants result from tuning some of the key-step mentioned above. There is an undeniable elegance in providing such a unifying view; even more important though is the fact that this unifying view also suggests novel ways of enhancing the performance of BP. We will also highlight how many contributions of the later chapters conform to this view as well.

3.5.2 Initialization

The choice of the initialization is no formality but, in case of multiple fixed points, determines which fixed point a given algorithm will converge to and as such it plays an important role in analyzing BPs behavior. Although only practical for problems of modest size, one can for example explore the solution space exhaustively by using a sufficiently wide range of initial values.

Note that BP often performs significantly better than worst-case analyses suggest. It seems, that for some reason, a wide range of initial message values converges to a good fixed point. This further raises the question of how one should initialize message passing algorithms. One common

choice, that often works well, is to initialize all messages to the same value $\mu_{ij}^1(x_j) = \frac{1}{|\mathcal{X}|}$; it remains unclear, however, where the advantages of uniform initialization stem from. In general, it would be desirable to identify initial message values that are optimal for certain problem-classes. We will later pursue this idea and provide a heuristic that obviates the need for choosing initial values and guides BP towards accurate fixed points (cf. self-guided belief propagation in Chapter 6).

3.5.3 Update

The choice of the update function is one of the key-aspects of message passing algorithms and has an enormous influence on the overall behavior of a particular algorithm. In particular, it is altering the update function that often leads to algorithms with favorable properties. The search for enhanced message passing algorithms thus spurred the development of alternative update functions, which ultimately led to the proposal of various algorithms. We will discuss some successful variants and group these modifications into two different categories.

First, instead of applying the update function \mathcal{F} to all messages at once, one can update only a single message at each iteration. Effectively, this procedure determines a schedule according to which the messages are updated; we consequently term this concept *scheduling*. Note that scheduling preserves all update functions and only changes the order in which they are applied.

On the other hand, one can, of course, try to improve the convergence properties by altering the update functions themselves. Although one can think of numerous ways to alter the update functions we intend to consider variants that perform at least as good as plain BP. In particular, this includes two variants that will be considered in this thesis; these are: *damping* and *self-guided belief propagation*.

Scheduling

It is widely accepted that scheduling the messages and updating them asynchronously enhances BP. In particular, this helps to achieve convergence more often and in fewer iterations. Up until now, we considered BP without scheduling where all messages are updated according to

$$\begin{aligned} \boldsymbol{\mu}^{n+1} &= \mathcal{BP}(\boldsymbol{\mu}^n) \\ &= (f_1(\boldsymbol{\mu}^n), \dots, f_m(\boldsymbol{\mu}^n), \dots, f_{2|\mathbf{E}|}(\boldsymbol{\mu}^n)). \end{aligned} \quad (3.43)$$

Note that we imposed some ordering $o : (i, j) \rightarrow m$ on the messages and refer to the update functions accordingly, i.e., f_m computes the m^{th} message. There are $2|\mathbf{E}|$ messages in total as they are passed along both directions of the edges.

For an asynchronous version of BP only a single message $\mu_{ij}(x_j) = \mu_m(x)$ is updated per iteration according to

$$\boldsymbol{\mu}^{(n+1)} = (\mu_1^n(x), \dots, f_m(\boldsymbol{\mu}^n), \dots, \mu_{2|\mathbf{E}|}^n(x)). \quad (3.44)$$

The update-order, i.e., which message to compute, has a major influence on the overall performance. One common choice is a *fixed order*, either according to a round-robin scheme or a random ordering. In both cases, the sorting of the messages follows a particular ordering and the messages are selected for an update according to

$$m = n \pmod{2|\mathbf{E}|}. \quad (3.45)$$

Such a fixed update-order already improves the performance of BP significantly, yet it leaves considerable room for improvement.

One popular way of scheduling the messages relies on the fact that, for a tree, BP is guaranteed to converge and to provide the exact marginals. *Tree based reparameterization* (TRP) [WJW02] computes a set of spanning trees so that the union of all spanning trees includes all edges of the original graph. Then, one applies BP to one tree at a time. One says that this calibrates the marginal along the spanning tree, i.e., the marginals would be exact if the graph would contain no more edges. Of course, performing BP on another spanning tree will alter some of the marginals again, but at least TRP enforces consistency through parts of the graph. The choice and the order of the spanning trees impact the overall performance of TRP and can be tailored in a model-specific way. Intuitively, one can argue that the focus on spanning trees enhances the overall convergence properties by enforcing a brisk exchange of information in the graph.

All scheduling methods discussed so far rely on a fixed update-order that may or may not depend on the graph. Instead of carefully selecting an update-order, it is, however, much more flexible to resort to an *adaptive* update-order, where the current message values determine which message to update next. This obviates the need for carefully tuning the update-order to the current graph and further improves the convergence properties because of the inherent flexibility in adaptive methods.

The most basic adaptive scheduling method, *residual belief propagation* (RBP) [EMK06], already improves the convergence properties a lot. RBP updates the messages along the least calibrated edge, i.e., for which the message varies the most. The underlying assumption is that strongly varying messages contain more information and are thus of predominant importance for the convergence of BP. To make this more precise we introduce the residual r_{ij} to measure the message-distances between two successive iterations¹⁹

$$\begin{aligned} r_{ij}^n &= \|\mu_{ij}^{n+1}(x_j) - \mu_{ij}^n(x_j)\|_\infty \\ &= \max_{x_j} |\mu_{ij}^{n+1}(x_j) - \mu_{ij}^n(x_j)|. \end{aligned} \quad (3.46)$$

Let $\mathbf{r}^n = \{r_{ij}^n : (i, j) \in \mathbf{E}\}$ be the set of all residuals at iteration n , then the index that maximizes the residual

$$m = \arg \max_m \mathbf{r}^n \quad (3.47)$$

identifies the message that is selected for an update. Overall this significantly reduces the number of message updates until convergence and the overall convergence time. Computing all residuals, however, poses some computational overhead. We can reduce this overhead and further reduce the convergence time without degrading the accuracy much by computing an upper bound on the residual [SM07].

Although adaptive scheduling increases the number of models for which BP converges, some models remain problematic. Closer inspection reveals that only small parts of the graph fail to converge, whereas large parts (almost) converge in a couple of iterations. RBP thus updates the same subset of messages ad infinitum. This observation inspired the design of scheduling methods that detect such problematic behavior, specifically counteract the oscillations, and enforce convergence [KRTP15].

Noise injection belief propagation (NIBP) applies RBP and actively checks for oscillations: if RBP converges no modifications are undertaken; if, however, oscillations occur, NIBP *injects* Gaussian noise to the selected message. The underlying intuition is that the injected noise

¹⁹ Ultimately one is interested in the distance to the fixed point, if it exists, $\lim_{n \rightarrow \infty} \mu_{ij}^n(x_j)$. However, since $\lim_{n \rightarrow \infty} \mu_{ij}^n(x_j)$ is not known, the time variation of the messages offers a valid surrogate (cf. [EMK06]).

will propagate from the most influential part of the graph (i.e., the message selected by RBP) throughout the whole graph and, by introducing a relevant change to the overall model, lead to convergence. Let $\mu_{ij}(x_j)$ be the message selected according to (3.47); if this message oscillates, i.e., if $\mu_{ij}^n(x_j) = \mu_{ij}^{n-l}(x_j)$ for any $l \in \mathbb{R}$, Gaussian noise is added so that

$$\mu_{ij}^{n+1}(x_j) = f_{ij}(\boldsymbol{\mu}^n) + \mathcal{N}(0, \sigma^2). \quad (3.48)$$

Alternatively, instead of actively detecting oscillations, one can question the underlying assumption of RBP. Messages, despite varying significantly, often take identical values repeatedly; arguably, these messages carry not much information. *Weight decay belief propagation* (WBP) penalizes this behavior and dampens the residual of messages that were already updated. Consequently, WBP increases the relevance of the remaining messages and thus further refines the parameterization of the overall graph. More precisely, WBP divides the residual of a messages by the number of times this message has already been scheduled and reformulates (3.47) according to

$$m = \arg \max_m \frac{r_{ij}^n}{\sum_{n'=1}^n \mathbf{1}_{ij}^{n'}}, \quad (3.49)$$

where the indicator function $\mathbf{1}_{ij}^{n'} = 1$ if and only if $r_{ij}^{n'} = \arg \max_m r_{ij}^{n'}$. Detailed pseudocode for NIBP and WBP is presented in Appendix A.3.




Figure 3.4: Percentage of converged runs % over the number of iterations n for: ABP (black), RBP (green), NIBP (blue), and WBP (orange). The results are averaged over 233 grid graphs of size 13×13 with random potentials specified by $(J_{ij}, \theta_i) \sim \mathcal{U}(-6.5, 6.5)$ in accordance with [EMK06].

When comparing the convergence properties of the discussed scheduling methods on grid graphs with random Ising potentials, a clear picture emerges. We present one exemplary experiment in Figure 3.4 and refer to the paper [KRTP15] for more exhaustive experiments.

Although scheduling with a round-robin schedule, ABP (black), improves the convergence properties over BP without scheduling, ABP does not perform as well as more sophisticated methods. In particular, adaptive scheduling enhances the convergence properties notably, demonstrating its conceptual advantage. Evaluating the overall convergence behavior, NIBP (orange) converges for more models than any other method. WBP (blue) still converges more often than RBP (green), although it requires more iterations to converge. Damping of the residual slows the convergence as messages that require multiple updates to converge will now be selected less frequently.

Although the positive influence of adaptive scheduling methods on the convergence properties is indisputable, the influence on the marginal accuracy is less obvious. In practice, however, we are not only interested in good convergence properties but require accurate marginals as well. A comprehensive comparison of different scheduling methods in terms of marginal accuracy is presented in [KRTP15]. We summarize the major insights here: considering only graphs for which a simple round-robin schedule (i.e., ABP) achieves convergence, ABP approximates the marginals well for all models, whereas RBP and NIBP fail to do so. It may seem that advanced scheduling methods pay the price for improving the convergence properties and obtain less accurate marginals. It is thus particularly remarkable that WBP maintains the marginal accuracy of ABP although it converges for a wide range of models. Of all considered methods, WBP is thus superior.

Alternative Update Functions

When the BP messages oscillate (often just between two values) applying a damping term has proven to be a successful strategy to improve the convergence properties [MWJ99]. That is, one replaces the messages with a weighted average of the last messages so that

$$\mathcal{BP}_D(\boldsymbol{\mu}) = (1 - \epsilon)\mathcal{BP}(\boldsymbol{\mu}) + \epsilon\boldsymbol{\mu}, \quad (3.50)$$

where $\epsilon \in [0, 1)$. Remember that we intend to change the update function without changing the fixed points. Complying with our requirement, fixed points of $\mathcal{BP}_D(\boldsymbol{\mu})$ must, therefore, be fixed points of BP without damping as well. This is the case as fixed points of BP must satisfy $\boldsymbol{\mu}^\circ = \mathcal{BP}(\boldsymbol{\mu}^\circ)$, which holds per definition so that

$$\mathcal{BP}_D(\boldsymbol{\mu}^\circ) = (1 - \epsilon)\mathcal{BP}(\boldsymbol{\mu}^\circ) + \epsilon\boldsymbol{\mu}^\circ = \mathcal{BP}(\boldsymbol{\mu}^\circ). \quad (3.51)$$

Note that, although damping does not affect the fixed points, the local stability of the individual fixed points may change (cf. Section 5.4.2).

In Chapter 6 we will propose self-guided belief propagation that relies on an alternative update function. While damping modifies the update function in a static manner, SBP modifies it in a more flexible, time-dependent, way. The underlying idea is to consider the potentials as a function of the iteration $\Psi(n)$, which enhances the performance across all areas (i.e., convergence properties and accuracy).

3.5.4 Evaluation

BP often performs better than suggested by worst-case analyses [MM09, Section 22.3]. This observation comes not too surprising, especially when considering the existence of multiple fixed points. We would indeed be rather surprised if all fixed points have similar accuracy. Rather than a disadvantage, the difference in the accuracy can be seen as a chance; modified evaluation functions may exploit this difference and select or combine fixed points to enhance the approximation quality. Certainly, modified evaluation functions are not a remedy to all problems. But modified evaluation functions are an important, flexible, concept to account for the existence of multiple fixed points and to mitigate some of the problems arising in loopy models.

Here we discuss two popular modifications that take advantage of the variability in accuracy across multiple fixed points. Taking the set of all fixed points \mathbf{T} (or a subset thereof) into account, we modify the evaluation function $\mathcal{D}_P : \{\boldsymbol{\mu}^\circ\} \rightarrow \tilde{P}_B$ in two ways. Either to select a single fixed point or to compute a weighted combination of multiple fixed points. Note that we will consider and evaluate both modifications in Chapter 5 and theoretically analyze them in

Chapter 7. The modifications of \mathcal{D}_P become relevant in the presence of multiple fixed points. BP may converge to any fixed point as it does, in general, not favor a particular one over the other. Even if all fixed points are available, however, it remains unclear how one should select the fixed point that approximates the marginals best.

Variational methods minimize their objective and aim to find the global minimum, although – for non-convex \mathcal{F}_B – this cannot be done efficiently. Given a set of fixed points, we should thus select the fixed point that minimizes \mathcal{F}_B (and maximizes \mathcal{Z}_B), which is the global minimum of \mathcal{F}_B . We denote the associated pseudomarginals by

$$\tilde{P}_B^{MAX} = \arg \max_{\tilde{P}_B \in \mathbf{T}} \mathcal{Z}_B(\tilde{P}_B) = \arg \min_{\mathbb{L}} \mathcal{F}_B(\tilde{P}_B). \quad (3.52)$$

Note that the complete set of fixed points \mathbf{T} is generally not available and the maximization can thus only be performed over the set of available fixed points.

Recently an interest emerged in methods that obtain and combine multiple fixed points [BYGRS12, BMZ05, SRF16]. The combination of all fixed points \mathbf{M} lies at the heart of the replica symmetry breaking (RSB) assumption. This assumption emerged in statistical physics from the study of spin glass models and describes how the exact marginals decompose into a weighted sum of different contributions, which turn out to correspond to minima of \mathcal{F}_B .²⁰ If \mathcal{F}_B is non-convex, the RSB theory describes the decomposition of the exact solution into a convex combination of marginals that are weighted by their associated partition function so that

$$P_{X_i}(x_i) = \frac{1}{\sum_m \mathcal{Z}_B^m} \sum_{m=1}^M \mathcal{Z}_B^m \tilde{P}_{X_i}^m(x_i). \quad (3.53)$$

This representation can be attributed to [MPV87] and, rather than a theorem, it is a set of postulates. One underlying assumption is that the system does exhibit multiple fixed points (unique fixed points would falsely imply exact marginals otherwise); an accessible introduction to the RSB theory and all underlying assumptions can be found in [MM09, Chapter 19]. Despite its non-rigorous flavor, (3.53) has been verified for a wide range of problems (e.g., random SAT problems and spin glasses). In particular, many state-of-the-art solvers for combinatorial problems rely on the RSB theory [RG15]. Further note that, even if the local polytope is a strict outer bound on the marginal polytope, a convex combination of $\tilde{P}_{X_i}^m(x_i)$ (that correspond to edge points of \mathbb{L} [YFW05]) may consequently end up on an edge point of \mathbb{M} (which it must if it is exact).

The RSB theory states how to form the exact solution from pseudomarginals that were obtained by efficient approximate inference methods. This seems to contradict all results on the computational complexity of exact inference. Closer inspection resolves this apparent contradiction: the number of fixed points may grow exponentially with the model-size so that obtaining all fixed points is just as problematic as computing the marginals exactly. Indeed, this is one of the reasons that limit the practical applicability of the RSB assumption.

For the special case of constrained satisfaction problems, survey propagation [BMZ05] is one efficient way to evaluate (3.53). The extension to more general models, however, remains somewhat elusive [SRF16, RG14].

Nonetheless, besides being relevant from a theoretical point of view, we will also identify certain models for which only relatively few fixed points exist, which substantiates the practical relevance of the RSB assumption as well.

²⁰ In physics one deals with the decomposition of the Gibbs measure (i.e., the joint distribution) into a weighted combination of Bethe measures (that correspond to BP fixed points).

3.6 Conclusion

This chapter introduced BP as a message passing algorithm for approximate inference. We further presented two alternative perspectives that hone our understanding of BP. Namely, we discussed (i) the variational interpretation that relates all fixed points of BP to stationary points of the Bethe free energy and (ii) the consideration as dynamical systems that suggest various ways to enhance BP. Both concepts arguably provided insights into the solution space of BP, explained the behavior of BP, and thus benefited our understanding of BP. Besides discussing the properties of BP from different perspectives, one major purpose of the current chapter was to provide an introduction for the second part of this thesis (i.e., Chapter 5 - 7). Our discussion is far from being exhaustive and left many questions open. Below we list a couple of open questions worth pursuing that will be addressed in the remainder of the thesis and that will provide interesting insights into the nature of BP.

One of the most pressing needs when considering the variational interpretation is a thorough analysis of the energy landscape of the Bethe free energy (specifically for problems where $\mathcal{F}_{\mathcal{B}}$ is non-convex). Since every local minimum serves as a potential fixed point it would be of great relevance to characterize the solution space and to obtain all possible fixed points.

While the knowledge of all fixed points is valuable in its own, this does not reveal all properties of relevance for BP. In particular, the energy landscape does not explain the convergence properties and it remains an open, albeit important, question under which conditions one can expect BP to converge.

Furthermore, we defined the set of all stable fixed points \mathbf{S} in Section 3.3.2. Whether this set is available is questionable as the full set of fixed points is generally not available. It is consequently an important subject to come up with message passing algorithms that tend to converge towards the most accurate fixed point.

Following the discussion in Section 3.4, it is crucial to assign some qualitative measurements to the fixed points after convergence of BP. This is particularly important as the exact solution is not available in practice. It remains an ongoing research topic, however, to determine useful accuracy-bounds that can be evaluated efficiently. Moreover, it is not obvious at all how different measurements, e.g, bounds on the marginal accuracy or the partition function, relate to each other.

4

Discrete-Time Dynamical Systems: Solution Space Analysis

*"Inside a broken clock
Splashing the wine
With all the rain dogs
Taxi, we'd rather walk..."*
– Tom Waits

This chapter provides a brief, high-level, introduction to (nonlinear) dynamical systems and further introduces tools from computational mathematics, required in the subsequent chapters. Already in Section 3.5, we have witnessed some of the benefits that result from considering BP as a nonlinear dynamical system. Framing BP as discrete time map allows us to draw from the rich history of dynamical systems (cf. [Tes12, Sch00]) and will be particularly beneficial in studying BP's convergence properties. Before we discuss how to analyze the convergence properties of a discrete-time map, a proper definition of the terminology is required. Subsequently, we express the essential steps in analyzing discrete-time maps, present multiple ways of doing so, and compare the individual advantages.

We begin this chapter with the definition of discrete time maps in Section 4.1. We then illustrate how the problem of computing the fixed points reduces to solving a polynomial system in Section 4.2, before we discuss and compare different approaches for solving polynomial systems in Section 4.3. Finally, we introduce the notion of stability and show how to assess the stability of a fixed point in Section 4.4.

4.1 Discrete Time Map

Let us first introduce the notation of a discrete time map now:

Definition 4.1.1 (Discrete Time Map). A discrete time map is defined by an n -dimensional state vector $\mathbf{x} \in \mathbb{R}^n$ and a function $\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that acts on the state vector according to $\mathbf{x}(n+1) = \mathcal{F}(\mathbf{x}(n))$, where $n \in \mathbb{Z}_+^*$ is the time index. Note that the function \mathcal{F} is a composition of n functions $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$. We will denote the value of the state vector at time n by $\mathbf{x}^n = \mathbf{x}(n)$ so that

$$\mathbf{x}^{n+1} = \mathcal{F}(\mathbf{x}^n). \quad (4.1)$$

A discrete-time map is further said to be *nonlinear* if some (or all) functions $f_i(\mathbf{x})$ are nonlinear.

One of the overarching goals in the field of dynamical systems is to understand how the system responds and evolves for given initial conditions in order to gain knowledge about the global behavior. In particular, this knowledge can be gained by first obtaining all fixed points of a map, for which the state vector \mathbf{x} remains unaffected by the application of (4.1). More formally,

a *fixed point* of the map \mathcal{F} is a state vector $\mathbf{x}^\circ \in \mathbb{R}^n$ that maps to itself so that

$$\mathbf{x}^\circ = \mathcal{F}(\mathbf{x}^\circ). \quad (4.2)$$

The problem of obtaining all fixed points is formally introduced in Section 4.2 before we compare different approaches in Section 4.3. The knowledge of all fixed points alone, however, does not provide sufficient knowledge about the global behavior of a system and is only part of the whole picture. According to the definition in (4.2), fixed points are not affected by (4.1), though it is not revealed if and under which conditions \mathbf{x} will converge to a specific fixed point. We are, however, interested in determining which fixed points can be obtained through the repetitive application of (4.1), i.e., whether a fixed point is stable (cf. Section 4.4).

4.2 Finding Fixed Points of a Discrete Time Map

Fixed points of a discrete-time map are state vectors that remain unaffected by the map; the fixed points consequently correspond to the solutions of the *fixed point equations* $\mathbf{x}^\circ - \mathcal{F}(\mathbf{x}^\circ) = 0$ per definition. Thus, computing the fixed points boils down to solving a system of polynomial equations, i.e., of nonlinear algebraic equations. However, albeit straightforward in principle, solving the fixed point equations is by no means straightforward to do in practice; in particular if confronted with a nonlinear system.

Systems of polynomial equations arise quite naturally in many engineering applications and it is a classical and important problem in computational mathematics to solve them. But, despite their relevance and the long research-history, it is still an active field of research. Today, solving systems of polynomial equations brings together many different branches of mathematics (such as topology, numerical mathematics, geometry, and algebra) that are joined under the term *algebraic geometry*. However, the body of the corresponding literature is vast and – as different fields play together – requires deep knowledge in many subjects. Moreover, the topic is unfortunately mainly studied from a theoretical point of view, which often prevents the application to practical problems and poses a notable hurdle for non-experts.

Although many approaches for solving polynomial systems seem promising at first, it is a daunting task to understand their properties and capabilities enough as to select an appropriate method, well-suited for the task at hand. Indeed, the ability to deal with large systems of polynomial equations was a major contribution and a prerequisite for many results that are obtained in Chapter 5. Although it is possible to read those sections and to understand the ultimate insights, this thesis should be as self-contained as possible. Therefore we give a brief (and far from complete) summary of the most prominent ways for solving polynomial systems subsequently. In particular, we will discuss numerical methods (Section 4.3.1), symbolical methods (Section 4.3.2), and numerical polynomial homotopy continuation methods (Section 4.3.3).

4.2.1 Systems of Polynomial Equations

Before we delve into the subtleties that come with the problem of solving systems of polynomial equations and before we compare some fundamentally different approaches, we have to provide some definitions first. Note that our notation is adapted from [CLO92], where a comprehensive overview is provided and the underlying concepts for solving polynomial systems are described in great detail.

Definition 4.2.1 (Polynomial). A polynomial f in the variables $\mathbf{x} = \{x_1, \dots, x_n\}$, and with coefficients a_1, \dots, a_K in \mathbb{C}^n , is a function of the form

$$f(x_1, \dots, x_n) = \sum_{k=0}^K a_k x_1^{k_1} \cdots x_n^{k_n}. \quad (4.3)$$

We say that the *monomial* $\mathbf{x}^k = \prod_{i=1}^n x_i^{k_i}$ is of degree $|k| = \sum_{i=1}^n k_i$ and that f is of degree $\max_{0 \leq k \leq K} |k|$.

One fundamental concept in commutative algebra is the notion of a polynomial ring $\mathbb{K}[x_1, \dots, x_n]$. It consists of all possible polynomials (constructed as in (4.3)) in the variables $\{x_1, \dots, x_n\}$ with coefficients $a_i \in \mathbb{K}$ (see [CLO92, Appendix A] for a more formal definition).

We will only work with polynomials that have their variables defined over the complex numbers, i.e., $\mathbf{x} \in \mathbb{C}^n$ in this thesis. Note that we are actually only interested in polynomials with real variables, i.e., with $\mathbf{x} \in \mathbb{R}^n$. Some methods do, however, require the definition over complex variables. Moreover, numerical inaccuracies may introduce a non-zero, albeit negligible, imaginary part to an originally real solution. Considering only $\mathbf{x} \in \mathbb{R}^n$ leads to the failure of obtaining these solutions. Also, considering the complex domain does no harm as we can later easily choose the real solutions.

Now let us consider a *system of polynomial equations* that we intend to solve:

Definition 4.2.2 (Polynomial System). A polynomial system, or set of polynomial equations, \mathbf{F} consists of s polynomials f_1, \dots, f_s in the variables $\{x_1, \dots, x_n\}$ and is a function of the form

$$\mathbf{F}(x_1, \dots, x_n) = \begin{cases} f_1(x_1, \dots, x_n) \\ \vdots \\ f_s(x_1, \dots, x_n). \end{cases} \quad (4.4)$$

We are interested in obtaining the set of solutions for which all equations equate to zero. This defines a *variety*, a classical object in algebraic geometry.

Definition 4.2.3 (Set of Solutions (Variety)). A variety, or the set of solutions, $\mathbf{V}(\mathbf{F}) \subset \mathbb{C}^n$, is the set of variables \mathbf{x} for which all equations equate to zero, i.e.,

$$\mathbf{V}(\mathbf{F}) = \{\mathbf{x} \in \mathbb{C}^n : f_i(\mathbf{x}) = 0 \text{ for all } f_i(\mathbf{x}) \in \mathbf{F}(\mathbf{x})\}. \quad (4.5)$$

Note that the variety is defined over the complex numbers as well; we further define the set of solutions over the real and the strictly positive real numbers $\mathbf{V}_{\mathbb{R}^+}(\mathbf{F}) \subset \mathbf{V}_{\mathbb{R}}(\mathbf{F}) \subset \mathbf{V}(\mathbf{F})$.

4.3 Solving Polynomial Systems: A Comparison of Methods

We are now well prepared to finally present how systems of polynomial equations are actually solved. A great variety of approaches have been developed over the years; we present some of the most well-known approaches and introduce the underlying concepts to highlight some of their advantages and disadvantages. It shall be noted that this is by no means an exhaustive overview. Instead of discussing all methods in full detail, which would be beyond the scope of this thesis, we rather aim to outline the underlying concepts for each method.

Every method comes with its specific properties, that have to be considered and should ideally be aligned with the requirements of a given problem; in particular, one should answer the following questions before selecting one particular method to solve the polynomial system (cf. [SW05, p.68]):

- Do we require the set of all solutions $\mathbf{V}(\mathbf{F})$, or just a single possible solution?
- What are the specifications of the problem under consideration (sparsity, problem size, etc.) and which methods are well-suited to exploit those properties?
- Is the method already available in some software-package?
- Are there any hyper-parameters that influence the efficiency of the method and how much experience and effort is required in adapting the parameters to the problem under consideration?

The subsequently presented methods belong to – and are the more prominent examples of – three fundamentally different approaches: numerical solvers, symbolic methods, and numerical polynomial homotopy continuation (NPHC) methods. We will evaluate and compare the different approaches with respect to the above-mentioned properties. Note that although the most relevant properties for the subsequent chapters will be pointed out, it would be presumptuous to speak of an all-encompassing list of advantages and disadvantages.

4.3.1 Numerical Methods

One basic and well-established method for solving systems of nonlinear equations is the Newton-Raphson method (often just called Newton’s method) which is an iterative solver that progressively refines an initial guess to reach a solution (see [CLO05, pp.30] for a brief introduction, or some classical books on numerical analysis, e.g. [BF11, Chapter 2.3]).

The underlying idea is to approximate the considered system by a first-order Taylor approximation in an initial guess, and to successively refine this guess by proceeding along a series of first-order Taylor approximations.

Newton’s method is a powerful tool that often converges quickly to the solution, particularly if one has a good initial guess. Theoretical results on the convergence rate of such iterative solver are relatively well established (cf. [BF11, Chapter 2.4]) and have two particularly relevant implications for our work: First, an initial guess, sufficiently close to a solution, is required or, otherwise, the iterative solver may diverge or even exhibit chaotic behavior. Second, many different initializations are required to obtain multiple solutions, and it consequently remains problematic to obtain the full set of solutions with these methods.

4.3.2 Symbolic Methods

From a completely different point of view, symbolic methods [CLO92, CLO05] (e.g. Gröbner basis method, Border basis method, Wu’s method, and method of sparse resultant) rely on symbolic manipulation of the polynomial system and successive elimination of variables to obtain a simpler but equivalent form. In a sense, these methods generalize the Gaussian elimination method from linear systems into the nonlinear settings.

Before we discuss the details of symbolic methods, we need to introduce some fundamental notions from commutative algebra. One particularly relevant concept is the notion of an ideal.

Definition 4.3.1 (Ideal). Consider a polynomial system $\mathbf{F}(x_1, \dots, x_n)$ in a polynomial ring $\mathbb{K}[x_1, \dots, x_n]$ and some other polynomial system $\mathbf{H}(x_1, \dots, x_n)$ in the same polynomial ring. The ideal $\langle f_1, \dots, f_s \rangle$ is closed under addition and multiplication and defines the set of all polynomial systems that satisfy

$$\langle f_1, \dots, f_s \rangle = \left\{ \sum_{i=1}^s h_i f_i : \mathbf{H} \in \mathbb{K}[x_1, \dots, x_n] \right\}. \quad (4.6)$$

We say that the ideal $\langle f_1, \dots, f_s \rangle$ is generated by the basis $\mathbf{F}(\mathbf{x})$.

The above definition of the ideal implies that any element of the ideal equates to zero if $\mathbf{F}(x_1, \dots, x_n) = 0$ (cf. [CLO92, pp.30]). This is an immediate consequence of (4.6) and it has important implications for obtaining the variety $\mathbf{V}(\mathbf{F})$. Not that, in fact, the Hilbert's basis theorem states that every ideal $\langle f_1, \dots, f_s \rangle$ is generated by a finite set of equations. This implies that the variety of the ideal

$$\mathbf{V}(\langle f_1, \dots, f_s \rangle) = \{ \mathbf{x} \in \mathbb{K}[\mathbf{x}] : f(\mathbf{x}) = 0 \text{ for all } f(\mathbf{x}) \in \langle f_1, \dots, f_s \rangle \}$$

must equal the variety of its generator, i.e., $\mathbf{V}(\mathbf{F}) = \mathbf{V}(\langle f_1, \dots, f_s \rangle)$ (cf. [CLO92, Proposition 9]). This has one particularly important consequence, namely that two polynomial systems $\mathbf{F}(\mathbf{x})$ and $\mathbf{G}(\mathbf{x})$ that generate the same ideal, i.e., for which $\langle f_1, \dots, f_s \rangle = \langle g_1, \dots, g_s \rangle$, will also have the same variety, i.e., $\mathbf{V}(\mathbf{F}) = \mathbf{V}(\mathbf{G})$.

Note that every ideal has not just one but many different generators; this is a crucial prerequisite for the concept of symbolic methods. Moreover, this allows us to draw an analogy to linear algebra: in this analogy, the ideal corresponds to some subspace (that is closed under addition and multiplication) that is spanned by some vectors (that correspond to the polynomials in the generator).

Symbolic methods take a fundamentally different approach to solving systems of polynomial equations $\mathbf{F}(\mathbf{x})$ than numerical methods. Whereas numerical methods aim to estimate the solution directly, symbolic methods focus on computing an alternative generator $\mathbf{G}(\mathbf{x})$ with “good” properties. The overarching aim is to construct $\mathbf{G}(\mathbf{x})$ so that both generators have the same ideal $\langle g_1, \dots, g_s \rangle = \langle f_1, \dots, f_s \rangle$ – and thus identical solutions, and to concurrently impose certain properties on $\mathbf{G}(\mathbf{x})$ that facilitate computing the set of solutions. In particular, we hope that $\mathbf{G}(\mathbf{x})$ consists of equations that have a low degree and are in an upper triangle form (to take up the analogy with linear algebra, the aim is to find a more intuitive description of the subspace, e.g., by orthonormal vectors).

Gröbner Basis

One particularly successful way to construct a good basis is to compute the (reduced) Gröbner basis. Together with the Buchberger algorithm that computes such a basis, Gröbner bases were introduced in the Thesis of Bruno Buchberger [Buc65, Buc06].

Gröbner bases tend to consist of equations of low degree that are – if possible - in upper triangle form. A formal and more comprehensive treatment of the Gröbner basis method can be found in many textbooks (e.g., in [CLO92]); here we shall be content with illustrating the underlying concepts exemplary.

Example 6 (Gröbner Basis for Solving System of Equations).

Let us solve a simple polynomial system $\mathbf{F}(\mathbf{x})$ with three variables in three equations by utilizing the Gröbner basis method. Therefore, consider

$$f_1(\mathbf{x}) = x_1^2 + x_2^2 - 1 \tag{4.7}$$

$$f_2(\mathbf{x}) = (x_1 - 1)^2 + x_2^2 - 1 \tag{4.8}$$

$$f_3(\mathbf{x}) = x_1^2 + x_2^2 + x_3^2 - 2. \tag{4.9}$$

We then compute the Gröbner basis using Wolfram Mathematica and obtain an alternative basis according to

$$g_1(\mathbf{x}) = x_3^2 - 1 \tag{4.10}$$

$$g_2(\mathbf{x}) = 4x_2^2 - 3 \tag{4.11}$$

$$g_3(\mathbf{x}) = 2x_1 - 1. \tag{4.12}$$

Note that all equations $g_i(\mathbf{x}) \in \mathbf{G}(\mathbf{x})$ depend only on a single variable and have low degree, thus satisfying all desired requirements. One can consequently obtain all four solutions from $\mathbf{G}(\mathbf{x})$ in a straight-forward manner; these are $(x_1, x_2, x_3) = (\frac{1}{2}, \pm \frac{\sqrt{3}}{2}, \pm 1)$.

The Gröbner basis method obviously obtains the full set of solutions (as compared to only a single one by numerical methods). Furthermore, methods that obtain the Gröbner basis are particularly attractive for solving systems of polynomial equations as many algorithms are available and included in many software packages.

The application of the Gröbner basis method to large-scale problems with many variables demands a problem-dependent adaption of the algorithm, however, that requires good knowledge of algorithmic details. Even more problematic is the fact that Gröbner bases are unstable under small changes in the coefficients, which limits its application to problems with rational coefficients. We exemplify this numerical instability in Example 7.

Example 7 (Instability of Gröbner Basis).

The Gröbner Basis is numerically unstable with respect to small changes in the parameters. Let us exemplify this by the following system of equation $\mathbf{F}(\mathbf{x})$ (inspired by [KR05, Example 6.4.1]): we consider two ellipses defined by the following two polynomials in two variables

$$f_1(\mathbf{x}) = \frac{1}{4}x_1^2 + x_2^2 - 1 \quad (4.13)$$

$$f_2(\mathbf{x}) = x_1^2 + \frac{1}{4}x_2^2 - 1. \quad (4.14)$$

Again, the Gröbner basis provides a well-behaved system of polynomial equations

$$g_1(\mathbf{x}) = 5x_1^2 - 4 \quad (4.15)$$

$$g_2(\mathbf{x}) = 5x_2^2 - 4, \quad (4.16)$$

that yields all four solutions $(x_1, x_2) = (\pm \frac{2}{\sqrt{5}}, \pm \frac{2}{\sqrt{5}})$ immediately.

Now, assume we rotate the ellipses slightly by changing the underlying system of equations to $\tilde{\mathbf{F}}(\mathbf{x})$ with

$$\tilde{f}_1(\mathbf{x}) = \frac{1}{4}x_1^2 + x_2^2 - 1 + 0.001x_1x_2 \quad (4.17)$$

$$\tilde{f}_2(\mathbf{x}) = x_1^2 + \frac{1}{4}x_2^2 - 1 + 0.001x_1x_2. \quad (4.18)$$

Note that this minor rotation shifts the four solutions change only slightly (cf. Figure 4.1); in fact, the solutions to $\tilde{\mathbf{F}}(\mathbf{x})$ reside within a radius of $5 \cdot 10^{-4}$ of the solutions to $\mathbf{F}(\mathbf{x})$. On the other hand, however, a dramatic change occurs in the corresponding Gröbner basis, both in the coefficients and in the monomials:

$$\tilde{g}_1(\mathbf{x}) = 1250x_2^3 - 1000x_2 + x_1 \quad (4.19)$$

$$\tilde{g}_2(\mathbf{x}) = x_2^4 - 1.6x_2^2 + 0.64. \quad (4.20)$$

Not only do the Gröbner bases $\mathbf{G}(\mathbf{x})$ and $\tilde{\mathbf{G}}(\mathbf{x})$ differ in their structure, but – even more problematic – computation of $\mathbf{V}(\tilde{\mathbf{G}})$ now suffers from numerical issues. The large coefficients in $\tilde{g}_1(\mathbf{x})$ introduce an extreme sensitivity with respect to numerical accuracy and thus render the computation of the solutions by backward substitution highly problematic.

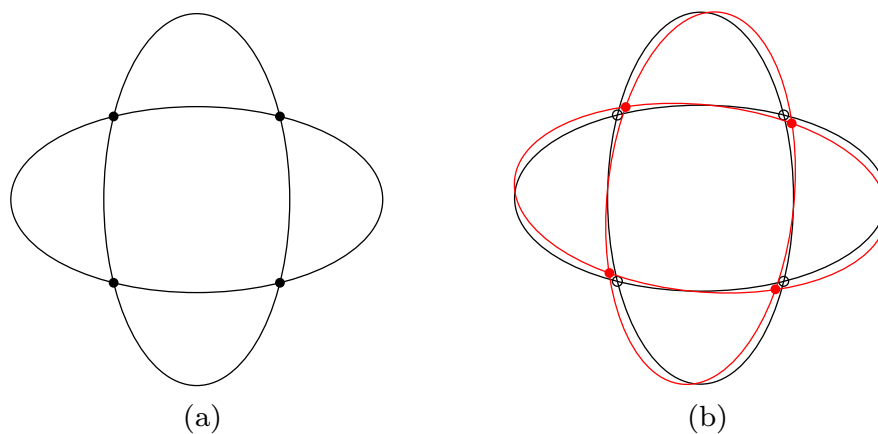


Figure 4.1: (a) Two ellipses defined by $F(\mathbf{x})$ (illustrated in black), where the solutions are depicted by solid black dots; (b) two rotated ellipses defined by $\tilde{F}(\mathbf{x})$ (illustrated in red), where the solutions are depicted by solid red dots (with the solutions of $F(\mathbf{x})$ depicted by black circles). Note that the solutions of both equation systems are similar.

According to the implicit function theorem, the solutions of $\tilde{F}(\mathbf{x})$ will change only slightly if the coefficients are modified by an infinitesimally small value; the Gröbner basis, however, may experience a dramatic shift. In fact, the result will also differ depending on the representation (decimal, or fraction of two integers) of the coefficients in $\tilde{F}(\mathbf{x})$. This poses a drastic problem for any polynomial system with non-rational coefficients as the representation of the coefficients will have a major influence on the final result.

Other Symbolic Methods

The limitation of the Gröbner basis approach to problems with rational coefficients and the numerical instabilities pose a considerable problem for many applications. Consequently, to extend the applicability of symbolic methods, much attention was focused on circumventing these issues and extending the range of solvable polynomial systems.

All symbolic methods adhere to one common characteristic; that is the search for an alternative well-behaved generator. There are various ways of obtaining this generator. One promising approach relates the computation of the generator (non-linear in the polynomial ring) to a similar problem in the quotient ring. The consideration in the quotient ring renders the problem a linear one that can be solved efficiently with well-established methods from linear algebra, e.g., by an eigendecomposition. This fruitful connection was first recognized in the work of Winfried Auzinger and Hans J. Stetter.²¹ Their work already revealed many expedient properties that come with working on the quotient ring [AS88, AS89]. This was only a first step, however, one that avoided the actual construction of the quotient ring [Ste04]. Nonetheless, the underlying ideas proved to be worth further pursuing and inspired a range of results. The inherent connection between the representation in the quotient ring and the problem of solving algebraic equations was first made explicit in [Möl93]. This connection finally led to the first working algorithm, the Möller-Stetter method [MS95]. An excellent overview of the underlying concepts and the initial developments can be found in [Ste04].

After laying out the foundation, many promising advances were proposed, often encompassed under the name of Border basis. This development is still ongoing and the advances are scattered over many publications. We refer to some of the excellent comprehensive papers available that provide a good overview of the recent advances and may serve as a good starting point for the

²¹ Stetter who was already about to retire found interest in the topic and prolonged his active research time by another 10 years; the results of which are summarized in his book [Ste04].

interested reader. The beneficial properties of Border basis are outlined and discussed in the tutorial paper [Mou07]; likewise [KKR05] summarizes the major properties and provides some accessible examples that show how to compute the Border basis. In a nutshell, the Border basis method combines the advantages of the Gröbner basis method, while avoiding some of the most relevant disadvantages (e.g., numerical instabilities). Unfortunately, however, much of the results remain of academic nature and have not yet found their way into existing software packages, thus severely limiting the potential application to practical problems.

We would like to stress that the toolbox of symbolic methods is of course not limited to Gröbner- and Border basis but contains many more methods, all with their own set of advantages and drawbacks. Nonetheless, we hope that this section provides a gentle introduction to some of the more established methods; a more detailed treatment of symbolic methods is available in a number of books (and the references therein) [CLO92, CLO05, KR00, KR05, DE05].

As a final statement, we conclude by stressing the overall drawback of symbolic methods: it is often expensive to apply these methods to problems of high degree. This is particularly critical as the sparsity in polynomial systems cannot be utilized in a straightforward manner so that even small – but high dimensional – systems suffer from this drawback.

4.3.3 Numerical Polynomial Homotopy Continuation (NPHC) Method

Another important approach for solving a system of polynomial equations is the *numerical polynomial homotopy continuation* (NPHC) method [Li03, SW05]. The NPHC method performs multiple stages in order to compute the solutions of the *target system* $\mathbf{F}(\mathbf{x})$.

First, the target system is inspected and a root count is computed that provides an upper bound on the number of solutions. Second, a closely related *start system* $\mathbf{Q}(\mathbf{x})$ is created that is trivial to solve and has precisely as many solutions as suggested by the root count. Third, the start system is continuously deformed into the target system. Finally, with appropriate construction, the trivial solutions of the start system also vary continuously under this deformation forming *solution paths* that connect to the desired solutions of the target system.

For instance, one basic form of a homotopy is given by

$$\mathbf{H}(\mathbf{x}, t) = (1 - t)\mathbf{Q}(\mathbf{x}) + \gamma t\mathbf{F}(\mathbf{x}) = 0, \quad (4.21)$$

for $t \in (0, 1]$ and with $\gamma \in \mathbb{C}$. Clearly, at $t = 0$ the homotopy reduces to the start system $\mathbf{Q}(\mathbf{x})$ and at $t = 1$ it reduces to the target system $\mathbf{F}(\mathbf{x})$. As t varies continuously from 0 to 1, the homotopy represents a deformation from the start system to the target system and the NPHC method tracks the solutions.

Example 8 (NPHC method for Solving an Equation).

Although the NPHC method is particularly suited for large systems of equations with many variables, we consider one minimalistic example for illustrative purposes. Let us consider the following target system consisting of a single equation $f(x)$ in a single complex variable x

$$\mathbf{F}(\mathbf{x}) = x^2 + ix - 2 = 0. \quad (4.22)$$

The root count for (4.22) is straightforward to estimate and just looking at the system reveals the existence of two solutions. This consequently requires a start systems with two solutions as well; one of the simplest system of equations with two solutions is for example given by

$$\mathbf{Q}(\mathbf{x}) = (x - 1)(x + 1) = 0. \quad (4.23)$$

We can immediately recognize both initial solutions of the start system $x_s = (-1, +1)$.

The homotopy $\mathbf{H}(x, t)$ is finally given by

$$\mathbf{H}(x, t) = (1 - t)(x - 1)(x + 1) + \gamma t(x^2 + ix - 2) = 0, \quad (4.24)$$

where $\gamma = \exp(i\theta)$ with a random angle $\theta \in [0, 2\pi)$. A couple of solution paths, that emerge for different values of θ , are illustrated in Figure 4.2. One can see that the angle θ influences the shape of the solution path. Every chosen value θ has a well-behaved solution path that connects the initial solutions of the start system $(x_{s,0}, x_{s,1})$ to the desired solutions of the target system $(x_{t,0}, x_{t,1})$. Using a predictor-corrector method to proceed along the solution path, the NPHC method then obtains both solutions of the target system $x_t = (\frac{\sqrt{7-i}}{2}, -\frac{\sqrt{7-i}}{2})$.




Figure 4.2: Solution paths for $\theta \in \{0, 1, \pi/2\}$. All solution paths lead from the solutions of the start system $x_{s,i}$ to the solutions of the target system $x_{t,i}$.

Although the NPHC method seemingly provides a simple method to obtain the desired solutions we have spared some of the subtleties that may arise; there are two ingredients in particular that require careful attention and were not mentioned in detail so far. These are the *path-tracking algorithm* that keeps track of the solution as t increases, the computation of the *root count*, and the creation of the *start system*.

Tracking the Solution

Because of its importance, the process of tracking the solution paths with increasing t is central to the NPHC method. It is thus important to consider robust and proficient path-tracking algorithms. We will present the key-steps of such path-tracking algorithms in the context of the NPHC method, closely following [SW05, Section 2.3], and refer to [AG03] for a more exhaustive description of path-tracking in general.

A generic approach must proceed along the solution paths, defined by the continuous homotopy $\mathbf{H}(\mathbf{x}, t)$ with known initial values $\mathbf{H}(\mathbf{x}, 0) = 0$, using a predictor-correction method. Such an iterative algorithm proceeds according to the following three steps. In the prediction step, one fits a function to the current values and predicts the values of \mathbf{x} at the next step. This can be done either by extrapolating a linear (or higher-order) function based on the last couple of points, or by linearizing in the last point and proceeding along the tangent direction. Note that the latter approach is well known as Euler's method.

In the correction step, one corrects the predicted value, using a numerical method as e.g., Newton's method. Despite its shortcomings discussed before, Newton's method is well suited for the correction step as the prediction step will estimate a value close enough for Newton's

method to converge. In general, the accuracy of the prediction will increase with a reduced step-size; if necessary, it is thus always possible to enforce convergence by adaptive reduction of the step-size.

Before repeating the above procedure and predicting the next value again, the step-size should ideally adapt according to the current correction step. If the preceding prediction was very accurate it is often safe to assume that the step-size can be enlarged, whereas it is preferable to reduce the step-size otherwise. Overall, the step-size significantly impacts the overall performance; a step-size too small introduces unnecessary many iterations and a step-size too large may lead to failure in the correction step.

These steps, prediction, correction, and step-size adaption are then repeated until the solutions of the target system are obtained for $t = 1$.

An important prerequisite for path-tracking algorithms is the existence of a well-behaved path. Fortunately, the introduction of some generic hyper-parameter, as for example the random complex number γ in (4.21), produces such well-behaved paths with probability one [SW05, Lemma 7.1.3].²²

Start System

Another crucial detail in the NPHC method is the computation of the root-count and the creation of the start system. There are various ways of computing this bound that differ in their complexity.

Loose bounds are often straightforward to compute but may lead to – potentially many – solutions in the start system that do not correspond to any solution in the target system. Still, one has to track all emerging solution paths, before finally recognizing the superfluous ones that diverge as the homotopy resembles the target system. Besides unnecessarily wasting computational resources this may have even more serious implications; a root count too large may render solving a problem infeasible because of the sheer amount of paths that need to be tracked.

Various methods have been proposed to bound the number of solutions (cf. [SW05, Section 8.1]) with more sophisticated methods generally reducing the number of paths to be tracked. Nonetheless, it is safe to assume that the computational effort required for estimating the bound increases with aiming for a tighter bound. Moreover, closely related to the estimation of the root count, the creation of the start system becomes more intricate as well.

For some problems, the consideration of more sophisticated methods scales nicely with the problem and the reduction in the number of paths outweighs the additional burden of creating an appropriate start-system. For other problems, however, estimating a more accurate root count requires significantly more computational resources than the second task, of tracking all paths to their solutions, altogether.

The optimal choice of method – that introduces as few paths as possible and spends as few resources for the creation of the start system as needed – is thus vital for the overall performance of the NPHC method. Unfortunately, general guidelines do not exist and choosing a reasonable method for a given problem is usually based on experience.

²² The main argument of the proof is that for all possible values of $\gamma \in \mathbb{C}$, except for a finite amount of distinct values, the solution paths are well-behaved. Picking the value γ at random will therefore result in a well-behaved path with probability one.

Properties of the NPHC method

The NPHC method takes the structure and the sparsity of the target system $F(\mathbf{x})$ into account. If the number of solutions is small, a well-chosen start system will also have only a few solutions thus implicitly exploiting the sparsity of the equation system.

The fundamental concept is that every solution of the target system has one corresponding solution residing in the start system. Looking at (4.21) it becomes apparent that, as t goes from 0 to 1, every solution path is completely independent of all others. This property admits parallel implementations that track all solution paths independently after the creation of the start system. The NPHC method is thus, in principle, suited for solving much larger systems than symbolic methods that cannot utilize parallelism in such a way.

Moreover, the NPHC method is implemented and available in a couple of software packages that often have interfaces with established software packages as, e.g., Maple, Matlab, or Python. Existing software packages, only naming a few, include PHCpack [Ver99], bertini [BHSW], and Hom4PS-3 [CLL14]. The existence of such established software packages makes the NPHC method attractive from a practical perspective and lowers the entrance hurdle.

One fundamental disadvantage of the NPHC method is the need for working in the complex domain, even though only the real solutions are of interest. Although the real solutions are contained in the set of complex solutions, working in the complex domain may lead to some problems: On the one hand, the NPHC method may fail if some of the solutions are not zero-dimensional, i.e., the solutions are continua instead of points. Restricting ourselves to the real line, these solutions could be zero-dimensional; yet the NPHC method will inevitably fail for its need to work in the complex domain. On the other hand, from a more practical perspective, the consideration of the complex field may increase the computational requirements significantly. This is most obvious if considering a problem with only a few real solutions but a huge amount of complex solutions that induce an equally huge amount of solution paths to be tracked. As it is not possible to determine the complex solutions upfront, all solution paths have to be tracked before neglecting most of them only after the termination of the NPHC method.

4.4 Stability Analysis of Fixed Points

As discussed at the beginning of this chapter, the knowledge of the fixed points alone is not sufficient in getting an adequate understanding of the considered map. All fixed points remain constant under repeated application of the considered map per definition, but may be different in their nature. To account for that difference we classify the fixed points as stable or unstable, corresponding to the long-term behavior of the map.

We start by getting some intuition behind the meaning of stable and unstable fixed points. Therefore, consider a dynamical system description of how a ball moves through a surface with valleys and peaks. By neglecting all kinds of forces that occur in real-world, except for gravity, we get a particularly simple model: along the slope, the ball will always accelerate downwards and decelerate upwards. Fixed points are consequently flat points on the surface with zero slope. Assume that ball sits on the top of a peak; then, moving the ball slightly will throw it out of equilibrium as it starts rolling towards another fixed point: we refer to the peak as an unstable fixed point. If, on the other hand, the ball resides at the bottom of a valley, the ball will return back to the fixed point after moving it slightly: we refer to the valley as a stable fixed point.

In analogy, the stability of a fixed point can be defined by the trajectory that the state vector describes after setting the state vector to some value close to a fixed point and whether repeated iteration of the map brings it back to the considered fixed point or not.

Let us make this notion of stability more precise now: A fixed point \mathbf{x}° is (*asymptotic*) *stable* if a neighborhood $U(\mathbf{x}^\circ)$ exists such that any $\mathbf{x}^1 \in U(\mathbf{x}^\circ)$ converges to \mathbf{x}° ; i.e., for every (small)

value of $\epsilon > 0$ there is a maximum number of iterations N so that \mathbf{x}^n is ϵ -close to \mathbf{x}° for $n > N$. A fixed point \mathbf{x}° is *unstable* if \mathbf{x} diverges from the fixed point for some $\mathbf{x}^1 \in U(\mathbf{x}^\circ)$.²³

4.4.1 Stability Analysis for a Nonlinear Discrete Time Map

We are now going to discuss a simple, yet powerful, method to assess the stability of a fixed point \mathbf{x}° under a specific discrete-time map $\mathcal{F}(\cdot)$. Therefore, we only require some tools from linear algebra. Note that we can express every linear map as a linear system of equations, and write it in matrix form so that $\mathcal{F}(\mathbf{x}) = \mathbf{B}\mathbf{x}$ where the i^{th} equation is given by

$$f_i(\mathbf{x}) = \sum_{k=0}^n a_{i,k} x_k. \quad (4.25)$$

This representation already suggests how to assess the stability of a fixed point. First of all, according to (4.2) a fixed point remains unaffected under multiplication by \mathbf{B} . Consequently, a fixed point \mathbf{x}° is stable if repeated application of \mathbf{B} brings the state vector $\mathbf{x} \neq \mathbf{x}^\circ$ to the fixed point; accordingly, we call a fixed stable whenever

$$\mathbf{x}^\circ = \mathbf{B}^n \mathbf{x} \quad (4.26)$$

for all values $n > N$ with $n, N \in \mathbb{Z}$.

Taking a closer look at (4.26) reveals how to assess the stability: the effect of repeated multiplication by \mathbf{B} directly relates to the eigenvalues [Str16, Chapter 6]. It is a fundamental property of linear algebra that a state vector converges to \mathbf{x}° if all eigenvalues λ_i have a magnitude strictly smaller than one. Accordingly, we conclude that the linear system $\mathcal{F}(\mathbf{x})$ exhibits a stable fixed point if its associated system-matrix \mathbf{B} has all eigenvalues $|\lambda_i| < 1$.

The above analysis crucially relies on the matrix representation of the discrete-time map. Since non-linear maps do not permit this matrix representation it is not directly possible to generalize the stability analysis. Despite this limitation, we may just approximate the nonlinear map $\mathcal{F}(\cdot)$ by a linear one in every fixed point and analyze the linearized system instead. Fortunately, the Hartman-Grobman theorem [Har60, Gro59] (see [Tes12, pp.264] for a proof) allows us to do just that. More precisely, the theorem permits the simplification of treating the nonlinear system as a linear one in its fixed points (at least in most cases).

A one-dimensional system is linearized by taking the derivative; a multivariate system is linearized by computing all first-order partial derivatives and collecting them in the Jacobian according to

$$\mathcal{F}'(\mathbf{x}) = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_s}{\partial x_1} & \cdots & \frac{\partial f_s}{\partial x_n} \end{bmatrix} \quad (4.27)$$

Then, in accordance with linear system theory, we inspect the eigenvalues of the Jacobian. Therefore, let us denote the spectrum of the Jacobian, i.e., the set of all eigenvalues by

$$\Lambda(\mathcal{F}'(\mathbf{x})) = \{\lambda_1, \dots, \lambda_n\}, \quad (4.28)$$

²³ Note that we do not discuss marginally stable fixed points, i.e., fixed points around which the state vector does neither converge nor diverge. Nonlinear systems require a more involved analysis as considered here if marginally stable fixed points are of relevance.

and the spectral radius, i.e., the maximum magnitude of all eigenvalue by

$$\rho(\mathcal{F}'(\mathbf{x})) = \max_{1 \leq i \leq n} |\lambda_i|. \quad (4.29)$$

This means that – whenever $\rho(\mathcal{F}'(\mathbf{x})) \neq 1$ – it is possible for nonlinear systems to infer the stability of a fixed point by looking at the spectral radius of the linearized system. A fixed point \mathbf{x}° is stable if all eigenvalues of the Jacobian have absolute value strictly smaller than one and lie inside the unit circle, i.e., if

$$\rho(\mathcal{F}'(\mathbf{x}^\circ)) < 1. \quad (4.30)$$

The system $\mathcal{F}(\cdot)$ consequently converges to \mathbf{x}° if initialized sufficiently close enough. A fixed point is unstable with respect to $\mathcal{F}(\cdot)$ if at least one eigenvalue exists outside the unit circle so that $\rho(\mathcal{F}'(\mathbf{x}^\circ)) > 1$.

The concept of linearization now allows us to assess the stability of a given fixed point. Note that this still excludes any discussion about the region of attraction; just because a fixed point is stable under a given map does not imply that every possible initialization will converge to the fixed point. Nonetheless, the notion of local stability has its merits and tells us whether a map can, in principle, converge to a fixed point (if it is stable) or if it can never converge to a fixed point (if it is unstable).

5

Solution Space of Belief Propagation: Number of Fixed Points and Their Stability

*” Travel makes one modest.
You see what a tiny place
you occupy in the world.”
– Gustave Flaubert*

In this chapter, we represent BP as a dynamical system and analyze its solution space in order to gain a deeper understanding of BP’s properties. Choosing from the variety of methods presented in Chapter 4 for solving the fixed point equations of BP, the NPHC method proves to be capable of providing the full set of fixed points. This is the first time that this is achieved for finite-size models with non-vanishing local potentials, which reveals a fundamental connection between the Bethe free energy and the accuracy of the marginals. The notion of a local stability analysis relies on the tools presented in Section 4.2 and was first applied to BP – although only for models without local potentials – in [MK05]. The NPHC method makes the full set of fixed points available and consequently admits an extension of the local stability analysis to more general models. This extension has far-reaching implications. For example, it was a common conjecture that strong local potentials positively influence the stability; we finally prove this assumption true. The empirical results further inspire the derivation of theoretical results that explain the influence of the potentials and the graph size on the convergence properties.

We begin this chapter with a brief summary of existing results on the number of fixed points in Section 5.2. In Section 5.3 we show how to formulate the fixed point equations for BP and how to use the NPHC method for solving them. We then reparameterize the messages and discuss the subtleties of applying a stability analysis to BP in Section 5.4. Finally, we specify a range of models in Section 5.5, before we compute all fixed points and evaluate their accuracy in Section 5.6 as well as their stability in Section 5.7. We then follow our empirical observations and analyze how the model-size and the potentials affect the stability in Section 5.8. In Section 5.9 we compute all fixed points and discuss the limitations of BP in the context of error-correcting codes.

Large portions of this chapter have been previously published but were considerably modified to fit into one coherent chapter: The computation of the full set of fixed points and the evaluation of the accuracy have been published in [KMCP18] and in [KPMC16]. Shifting the focus away from symbolic methods was triggered by a fruitful discussion with Michael Kerber; the application of polyhedral homotopy methods finally proved to be successful with the aid of Dhagash Metha and Tianran Chan. In particular the interpretation of the results and the assessment of the accuracy of the BP fixed points are due to the present author. The local stability analysis is a result of joint work with Franz Pernkopf and is published in [KP17].

5.1 Motivation

The previous chapters introduced BP as an *efficient* method for approximate inference. For arbitrary models with many loops, however, neither guarantees for convergence nor bounds on

the approximation error are established. It is precisely the existence of many loops, however, that renders exact inference intractable and requires the utilization of approximate inference methods (as for example BP). A better understanding of BP’s capabilities and limitations would therefore be of great relevance for loopy models. In this chapter, we aim to analyze BP on a range of models in great detail; we further aim to establish theoretical properties on the performance of BP and explain why BP often works surprisingly well on dense and loopy graphs, but other times fails to converge or give accurate results.

If multiple fixed points exist, the expected behavior of BP depends on the number of fixed points and the individual fixed point’s behavior. It is thus quite natural to strive for answers to the following questions:

- Can we specify model classes – by restricting, either individually or jointly, the structure and the parameters of probabilistic graphical models – for which a unique fixed point exists?
- Which models admit at least one fixed point to which BP converges for any choice of initial message values (or just if initialized close enough)? Does uniqueness of a fixed point imply that BP converges?

These questions have obviously been addressed before: sufficient conditions for uniqueness of fixed points were proposed by accounting for both the potentials as well as the graph structure [Hes04, MK07]. A direct relationship exists between accuracy and convergence rate for graphs with a single loop [Wei00] and for small grid graphs [Ihl07]. In contrast, graphs do exist that feature surprisingly accurate fixed points although BP fails to converge [WJ14a]. We shall thus gain a better understanding of the behavior of BP by developing a precise relation among the *number of fixed points* (and conditions for uniqueness), the *approximation accuracy*, and the *convergence properties*.

The consideration of BP as a discrete-time map suggests one way to compute the solution space of BP; the analysis of which reveals deep insights into the behavior of BP and (at least partially) answers the questions raised above. Applying the tools prepared in Chapter 4, one has to perform the following three steps for a given model: (i) obtain the set of all possible fixed points; (ii) assess and compare the accuracy for all fixed points; and (iii) analyze the stability of all fixed points.

First, we must obtain all fixed points: BP is unsuited for this task of obtaining *all* fixed points as it provides only a *single* fixed point and, obviously, does not provide unstable fixed points. If the unstable fixed points correspond to local minima of the Bethe free energy they can be obtained by methods that minimize $\mathcal{F}_{\mathcal{B}}$ directly (cf. Section 3.3.3). Nonetheless, those methods fail to obtain fixed points that correspond to local maxima and are not even guaranteed to obtain all local minima. In order to find all fixed points we reformulate the fixed point equations as a system of polynomial equations that we solve directly.

Second, after all fixed points are obtained we assess their accuracy in a straightforward manner; this is possible since we keep the models small enough to admit the application of exact inference methods for comparison.

Third, the concept of local stability is well established for discrete-time maps in general, as well as for BP in particular (see [MK05] for the special case of binary pairwise models without local potentials). The major obstacle that prohibits the stability analysis for more general models is the need for all fixed points, which are not known in general. We have already obtained all fixed points in the first step, so that we can analyze the stability of all fixed points by linearization of BP now.

Finally, we consolidate the insights of all three steps and shed some light onto the relation between the number of fixed points, the approximation accuracy, and the convergence properties. In particular, we will provide novel insights and find answers to the following questions:

- Convergence properties and accuracy relate to each other for small models. Does this generalize to models of arbitrary size as well?
- Does BP favor a particular fixed point if multiple fixed points are present? Is the most accurate fixed point always stable?
- What can be said about the accuracy if multiple fixed points exist? Do accurate fixed points still exist or does the existence of multiple fixed points imply failure of BP to provide accurate results? Can we enhance the approximation quality by a (suitable) combination of multiple fixed points?
- Under which circumstances does damping help to enforce convergence of BP?
- What is the influence of the model parameters: How does the graph structure (number of variables and connectivity) influence the number of fixed points and their stability for finite-size graphs? How do the local potentials influence the accuracy and the stability of fixed points? – It is a common conjecture (cf. [MK05]) that models with weak local potentials perform worst; why does the presence of strong local potentials enhance the performance of BP?

5.2 Fixed Points of Belief Propagation

Let us recall the iterative message update equations as discussed in Section 3.1:

$$\mu_{ij}^{n+1}(x_j) = \alpha_{ij}^n \sum_{x_i \in \mathcal{X}} \Phi_{X_i, X_j}(x_i, x_j) \Phi_{X_i}(x_i) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(x_i). \quad (5.1)$$

The collection of all update equations $\boldsymbol{\mu}^{n+1} = \mathcal{BP}(\boldsymbol{\mu}^n)$ defines a discrete-time map. We can thus rely on the standard recipe for analyzing discrete-time maps and analyze the solution space induced by (5.1) as discussed in Chapter 4. We begin with computing and counting the fixed points $\boldsymbol{\mu}^\circ$ for some reasonable form of normalization.

One alternative way to compute and count the fixed points is to utilize the correspondence to stationary points of the Bethe free energy and to study this energy landscape instead. Several methods (cf. Section 3.3.3) are available that minimize $\mathcal{F}_{\mathcal{B}}$ and, in the presence of multiple BP fixed points (i.e., for non-convex $\mathcal{F}_{\mathcal{B}}$), either converge to a local or the global minimum. None of these methods, however, is guaranteed to capture all fixed points. Moreover, every method that relies on minimizing $\mathcal{F}_{\mathcal{B}}$ inevitably fails to account for (unstable) fixed points that correspond to local maxima of $\mathcal{F}_{\mathcal{B}}$.²⁴

5.2.1 Number of Fixed Points on Regular Ising Graphs

Remember that the Ising model exhibits critical regions in the parameter space (so-called phase transitions) where the behavior of the model changes abruptly, unless it is specified on a path graph [Geo11, Chapter 12]. In the physics literature, one computes these phase transitions for infinite-size, or at least very large, models. Smaller graphs render the computation of the

²⁴ The trivial (paramagnetic) solution $\mu_{ij}^\circ(x_j) = \frac{1}{|\mathcal{X}|} = \frac{1}{2}$ for all $\mu_{ij}(x_j) \in \boldsymbol{\mu}$ is known in the special case of models with $\theta = 0$ and, depending on the coupling strength, it is either a local minimum or a local maximum of $\mathcal{F}_{\mathcal{B}}$ [MK05]. The extension to models with $\theta_i \neq 0$, however, is not straightforward and the complete set of all stationary points is not known in general.

phase transitions much more intricate.²⁵ If, however, all nodes have equal degree (and unitary potentials $J_{ij} = J$ and $\theta_i = \theta$), the phase transitions can be computed for small graphs as well. Therefore, one replaces the finite-size graph by an infinite-size graph with identical properties, a so-called Cayley tree.²⁶ Then, the Cayley tree provides a way to compute the phase transitions of the underlying graph in an analytical fashion [TM06].

Let us partition the parameter space (θ, J) into three distinct regions (F) , (P) , and (AF) . This terminology complies with the naming convention in statistical physics where the regions are referred to as ferromagnetic (F) , paramagnetic (P) , and antiferromagnetic (AF) . We can then, in accordance with [Geo11, Section 12.2], define the phase transitions partitioning the parameter space. Therefore, let us first introduce the shorthand notation $w = \tanh |J|$; we further introduce the following function

$$p(J, d) = \begin{cases} d \operatorname{atanh} \sqrt{\frac{d-w-1}{d/w-1}} - \operatorname{atanh} \sqrt{\frac{d-1/w}{d-w}} & \text{if } J > \operatorname{arccoth}(d) \\ d \operatorname{atanh} \sqrt{\frac{d-w-1}{d/w-1}} + \operatorname{atanh} \sqrt{\frac{d-1/w}{d-w}} & \text{if } J < \operatorname{arccoth}(d) \\ 0 & \text{else.} \end{cases} \quad (5.2)$$

Finally, the regions are specified according to

$$(J, \theta) \in (F) \quad \text{if } J > 0, J > \operatorname{arccoth}(d) \quad \text{and } |\theta| \leq p(J, d), \quad (5.3)$$

$$(J, \theta) \in (AF) \quad \text{if } J < 0, J < -\operatorname{arccoth}(d) \quad \text{and } |\theta| < p(J, d), \quad (5.4)$$

$$(J, \theta) \in (P) \quad \text{if } (J, \theta) \notin (F) \quad \text{and } (J, \theta) \notin (AF). \quad (5.5)$$

For attractive models with positive couplings ($J > 0$) BP converges to a unique fixed point inside (P) . This fixed point becomes unstable and two additional fixed points emerge inside (F) [YFW05, MM09]. For repulsive models with negative couplings ($J < 0$) BP only converges inside (P) and not inside (AF) [MK05].

The phase transitions of the complete graph will be computed according to (5.3)-(5.5). The phase transitions of the grid graph can only be numerically estimated and are defined by sudden changes in the number of fixed points. We will later discuss the specific influence of the graph-size on the phase transitions in Section 5.8.

5.3 Fixed Point Equations of Belief Propagation

Now, let us consider the update equations (5.1) as a nonlinear discrete-time map (cf. Section 3.5). We will adhere to the notation of Chapter 4 and Definition 4.1.1 in particular and consequently formulate the fixed point equations for the update rule of BP. Then, we aim to solve the fixed point equations directly; this will yield the set of all BP fixed points.

We did not consider a particular form of message normalization so far, i.e., how to choose α_{ij} . The normalization, however, affects the update rule and must therefore be defined explicitly before attempting to solve the fixed point equations; we define the normalization terms α_{ij} so that the messages along every edge sum up to one, i.e., $\alpha_{ij}^n = \sum_{x_j \in \mathcal{X}} \mu_{ij}^n(x_j) - 1$.

Let us recall the definition of fixed point messages from (4.2). Then, the update equations

²⁵ Note that true phase transitions, for which the partial derivatives of $\mathcal{F}_{\mathcal{B}}$ vanish can only exist for infinite-size graphs [Bin87]; with slight abuse of notation we refer to the finite-size manifestations of phase transitions as phase transitions as well.

²⁶ A Cayley tree is an infinite tree without loops that captures the interactions of a cyclic finite-size graph.

define the following set of polynomial equations $\mathbf{F}(\boldsymbol{\mu}, \boldsymbol{\alpha}) = \boldsymbol{\mu}^\circ - \mathcal{BP}(\boldsymbol{\mu}^\circ) = 0$ so that

$$\mathbf{F}(\boldsymbol{\mu}, \boldsymbol{\alpha}) = \begin{cases} \mu_{ij}^n(X_j = +1) - \alpha_{ij}^n \sum_{x_i \in \mathcal{X}} \Phi_{X_i, X_j}(x_i, +1) \Phi_{X_i}(x_i) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(x_i) \\ \mu_{ij}^n(X_j = -1) - \alpha_{ij}^n \sum_{x_i \in \mathcal{X}} \Phi_{X_i, X_j}(x_i, -1) \Phi_{X_i}(x_i) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(x_i) \\ \mu_{ij}^n(X_j = +1) + \mu_{ij}^n(X_j = -1) - 1. \end{cases} \quad (5.6)$$

This system of polynomial equations consist of S equations $(f_1(\boldsymbol{\mu}, \boldsymbol{\alpha}), \dots, f_s(\boldsymbol{\mu}, \boldsymbol{\alpha}))$, where

$$S = 2|\mathbf{E}| \cdot (|\mathcal{X}| + 1). \quad (5.7)$$

It is advantageous to consider this polynomial system defined over the complex numbers instead of the real numbers. Else, we would restrict ourselves and rule out several methods for solving systems of polynomial equations that rely on the definition over the complex numbers (cf. Section 4.3.3). Let us now define the set of solutions over the complex numbers, without accounting for multiplicity, according to

$$\mathbf{V}(\mathbf{F}) = \{(\boldsymbol{\mu}, \boldsymbol{\alpha}) \in \mathbb{C} : f_i(\boldsymbol{\mu}, \boldsymbol{\alpha}) = 0 \text{ for all } f_i \in \mathbf{F}\}. \quad (5.8)$$

We are particularly interested in the set of solutions over strictly positive real numbers $\mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F}) \subseteq \mathbf{V}(\mathbf{F})$. Note that $\mathbf{V}_{\mathbb{R}_+^*}$ is of specific relevance as it directly relates to the fixed points of BP.

Theorem 1 (Fixed Points of BP). *Let $(\boldsymbol{\mu}, \boldsymbol{\alpha})$ be some set of messages and normalization terms. Then, $(\boldsymbol{\mu}, \boldsymbol{\alpha})$ is a fixed point of BP, if and only if $(\boldsymbol{\mu}, \boldsymbol{\alpha}) \in \mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$.*

Proof. First, we show that every $(\boldsymbol{\mu}, \boldsymbol{\alpha}) \in \mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$ characterizes a fixed point of BP. All messages are positive by definition and are normalized according to (5.6) so that they represent probabilities (cf. Lemma 1). Furthermore, it follows from (5.6) that $\boldsymbol{\mu} - \mathcal{BP}(\boldsymbol{\mu}) = 0$, which constitutes a fixed point.

Conversely, consider some fixed point messages with its corresponding normalization coefficients $(\boldsymbol{\mu}^\circ, \boldsymbol{\alpha}^\circ)$, it then follows by definition that $\mathcal{BP}(\boldsymbol{\mu}^\circ) = \boldsymbol{\mu}^\circ$ and consequently $\mathbf{F}(\boldsymbol{\mu}, \boldsymbol{\alpha}) = 0$. \square

The set of solutions over the strictly positive real number thus directly corresponds to the set of all fixed points according to $\mathbf{T} = \{\tilde{P}_B = \mathcal{D}_P(\boldsymbol{\mu}^\circ), \mathcal{Z}_B = \mathcal{D}_Z(\boldsymbol{\mu}^\circ) : (\boldsymbol{\mu}^\circ, \boldsymbol{\alpha}^\circ) \in \mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})\}$.

Corollary 1.1. *Consider a graph with strictly positive potentials $\Phi(x_i, x_j)$ and $\Phi(x_i)$ as e.g., with Ising potentials. Then, the solution set $\mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$ is nonempty.*

Proof. For non-negative potentials the average energy is bounded from below and $\mathcal{F}_{\mathcal{B}}$ has at least one minimum [YFW05, Theorem 4]. Minima of the constrained $\mathcal{F}_{\mathcal{B}}$ correspond to BP fixed point solutions, the existence of which implies non-emptiness of $\mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$ by Theorem 1. \square

5.3.1 Solving the Fixed Point Equations

Solving systems of nonlinear polynomial equations is a classical problem in computational mathematics and a great variety of methods have been developed such as iterative solvers (cf. Section 4.3.1), symbolic methods (cf. Section 4.3.2), and homotopy methods (cf. Section 4.3.3). We now briefly review the applicability of these methods for the purpose of computing the set of all solutions.

The main disadvantage of iterative solvers is the need for an already known initial guess in the vicinity of a solution. Moreover, it is difficult, to obtain the full set of solutions with these

methods; this is of particular relevance as we are specifically interested in the entire set of the (positive) real solutions $\mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$.

Symbolic methods, on the other hand, are capable of obtaining the entire solution set. In particular, the Gröbner basis method is available in a number of software-packages and has seen substantial development in the past several decades. As discussed in Section 4.3.2, the main drawbacks are: the limitation to rational coefficients (cf. Example 7), a worst-case complexity that is double exponential in the number of variables [MM82, MM84], and the limited scalability in parallel computations. Altogether, this limits the application of symbolic methods to smaller systems; in fact the Gröbner basis method did not converge for any of our considered models.

The NPHC method obtains the entire solution set by deforming a *start* system that is trivial to solve to the *target* system that we intend to solve. Although maybe not as established as the Gröbner basis method, multiple software packages are available for the NPHC method as well. One crucial ingredient of the NPHC method is that a number of independent solution paths is tracked under the deformation of the homotopy. This suggests to track all solution paths independently, which makes the approach *pleasantly parallelizable*; this is essential in dealing with large polynomial systems. Even though only positive real solutions are of interest in this work, the NPHC method requires us to extend the domain to the field of the complex numbers in order to guarantee the emergence of smooth solution paths. Obviously a notable computational overhead is introduced in doing so; especially if only few of the solutions fall onto the real line. Various different forms of homotopies are available that differ in the complexity of bounding the number of solutions. As discussed in Section 4.3.3 the optimal choice depends strongly on the particular target system and is rarely known up-front.

At large, the NPHC method is the most promising approach for a couple of reasons: it obtains *all* isolated nonzero complex solutions²⁷ that must include *all* BP fixed points $\mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$; it is readily available in software packages; and it has a level of parallel scalability that can deal with models of relevant size

5.3.2 Polyhedral Homotopy Method

Let us briefly recall one of the most basic forms of homotopies from Section 4.3.3:

$$\mathbf{H}(\mathbf{x}, t) = (1 - t)\mathbf{Q}(\mathbf{x}) + \gamma t\mathbf{F}(\mathbf{x}) = 0. \quad (5.9)$$

Alternatively, one can construct more advanced "nonlinear" homotopies where the parameter t appears in nonlinear form, in order to reduce the computational costs. Among a great variety of polynomial homotopy constructions, the *polyhedral homotopy method*, developed by B. Huber and B. Sturmfels [HS95], is particularly suited for the systems of polynomial equations considered.

In applying the NPHC method to solve (5.6), the choice of $\mathbf{Q}(\boldsymbol{\mu}, \boldsymbol{\alpha})$ (the trivial system of equations that the target system is deformed into) plays an important role in the overall efficiency of the approach since different choices of $\mathbf{Q}(\boldsymbol{\mu}, \boldsymbol{\alpha})$ may induce a vastly different number of solution paths one has to track. The crucial part is to come up with a good upper bound on the number of solutions and to create an appropriate start system. Once this is solved, the desired solutions are simply obtained by tracking all independent solution paths.

Note that in each equation of (5.6) only few of the monomials are present, i.e., the update equations of BP imply a sparse system of equations [HS95]. In our experiments, we observed that despite the rather high *total degree*²⁸ d_t [SW05, pp.118], each equation in (5.6) contains only

²⁷ Here, "nonzero complex solutions" refer to complex solutions of a system of polynomial equations where each variable is nonzero. A solution is considered to be isolated if it has no degree of freedom, i.e., there is an open set containing it but no other solutions.

²⁸ The total degree of a system of polynomial equations is the product of the degrees of each equation. It is a basic fact in algebraic geometry that the total number of isolated complex solutions a polynomial system has

relatively few of the monomials. Such sparse systems usually benefit from resorting to more involved method homotopy methods that take the structure of the system into account and consequently provide tighter bounds on the number of solutions. The number of solution paths one has to track when using the polyhedral homotopy method for solving a system of polynomial equations is given by the so-called *Bernstein-Kushnirenko-Khovanskii (BKK) bound*: fixing the list of monomials that appear in the polynomial system, it is an important yet surprising fact in algebraic geometry that for almost all choices of the coefficients (in the probabilistic sense), the number of isolated nonzero complex solutions is a fixed number which only depends on the list of monomials. This number is known as the BKK bound [Ber75, Kus76, Kho78]. Intermediate steps in the determination of the BKK bound are reused to create an appropriate start system. Using the fully parallel implementation Hom4PS-3 [CLL14] of the polyhedral homotopy method, we compute the BKK bound, that is tight in all our experiments, and obtain *all* isolated positive solutions.

The polyhedral homotopy method exploits the structure of (5.6), but also requires some subtle steps. Rather than presenting all technical details we present an illustrative example to explain the underlying principles. For more details we refer the reader to the excellent overview papers [Li97, Li03, CL15] or to [SW05, Section 8.5.4] and the references therein.

Example 9 (Polyhedral Homotopy).

The essential steps in solving polynomial system with the polyhedral homotopy method are: first, to compute a root count based on mixed volume computations [Li03, Section 3]; second, to come up with an easy to solve start system [Li03, Section 4]; and finally, to solve the start system and track the solution paths to the target system [Li03, Section 1].

(i) *Root Count*: Consider the example system (taken from [SW05, p.142]) with two unknown variables $\mathbf{x} = \{x_1, x_2\}$ and with 4 solutions.

$$\mathbf{F}(\mathbf{x}) = \begin{cases} 1 + ax_1 + bx_1^2x_2^2 \\ 1 + cx_1 + dx_2 + ex_1x_2^2. \end{cases} \quad (5.10)$$

The total degree of this system of equations is $d_t = 4 \cdot 3 = 12$, which serves as an upper bound on the actual number of solutions. If the system of equations is sparse, the BKK bound serves as a much tighter bound.

Every equation $f_i \in \mathbf{F}(\mathbf{x})$ has an associated polytope S_i which is the convex hull of the exponent vectors for all monomials of f_i . For f_1 the polytope is

$$S_1 = \{(0,0)(1,0)(2,2)\}, \quad (5.11)$$

which has a graphical representation in Figure 5.1 (a). Similar for f_2 the polytope, shown in Figure 5.1 (b), is

$$S_2 = \{(0,0)(1,0)(0,1)(1,2)\}. \quad (5.12)$$

Some important operation on polytopes are the computation of the Minkowski sum $S_1 + S_2 = \{s_1 + s_2 : s_1 \in S_1, s_2 \in S_2\}$ and the computation of volumes, denoted by $V(S_i)$. Note that computing the BKK bound is a viable thing to do in any dimension; therefore we refer to $V(S_i)$ as volume although the polytopes of this example only lie in the two-dimensional space. The computation of the mixed volume $M(S_1, S_2)$ is a combinatorial problem that is especially comprehensible in the case of two equations where

$$M(S_1, S_2) = V(S_1 + S_2) - V(S_1) - V(S_2). \quad (5.13)$$

is bounded by its total degree (i.e., Bezout bound). Therefore the total degree serves as a crude measure of the complexity of the polynomial system.

For a generalization to higher dimensions see [SW05, p.140]. The polytope of $S_1 + S_2$ is illustrated in Figure 5.1 (c); the mixed volume is accordingly obtained by subtracting $V(S_1)$ and $V(S_2)$ from $V(S_1 + S_2)$, which equals the sum of all gray areas (known as mixed cells).

It is straightforward to see that each parallelogram has volume equal to 2; the BKK bound therefore equals $M(S_1, S_2) = 4$ and thus provides a tight bound on the number of solutions.




Figure 5.1: (a) polytope S_1 , (b) polytope S_2 , and (c) Minkowski sum $S_1 + S_2$.

(ii) *Start System:* The BKK bound does provide a tight bound on the number of solutions but does not immediately reveal the initial solutions of an appropriate start system $\mathbf{Q}(\mathbf{x})$ with $q_i(\mathbf{x}) = \sum_{a \in S_i} c_{i,a} \mathbf{x}^a$, where $\mathbf{x}^a = x_1^{a_1} \cdot x_2^{a_2}$ and $c_{i,a}$ are random coefficients.

However, the mixed volume computation can also be accomplished by introducing a lifting

$$\omega_i = \{\omega_i(a) : a \in S_i\}$$

for each f_i . Thereby, we increase the dimension of the polytope S_i to \hat{S}_i by adding one component to each exponent-vector a . This component is obtained by the lifting function $\omega_i(a)$. In our example we choose the lifting values $\omega_1 = \{0, 0, 0\}$ and $\omega_2 = \{0, 1, 1, 3\}$; these values are obtained by the inner products $\omega_1(a) = (0, 0) \circ (a_1, a_2)$ and $\omega_2(a) = (1, 1) \circ (a_1, a_2)$. The polytopes are lifted accordingly so that

$$\begin{aligned} \hat{S}_1 &= \{(0, 0, 0)(0, 1, 0)(2, 2, 0)\}, \\ \hat{S}_2 &= \{(0, 0, 0)(1, 0, 1)(0, 1, 1)(1, 2, 3)\}, \\ \hat{S}_1 + \hat{S}_2 &= \{(0, 0, 0)(0, 1, 0)(2, 2, 0)(0, 2, 1)(0, 1, 1)(3, 2, 1)(3, 4, 3)\}. \end{aligned}$$

Then the faces in the lower hull of $\hat{S}_1 + \hat{S}_2$ correspond to cells shown in Figure 5.1, which is known as a fine mixed subdivision.

These liftings, together with the random coefficients $c_{i,a}$, now form the homotopy $\hat{\mathbf{Q}}(\mathbf{x}, t)$ with $\hat{q}_i = \sum_{a \in S_i} c_{i,a} \mathbf{x}^a t^{\omega_i(a)}$ such that

$$\hat{\mathbf{Q}}(\mathbf{x}, t) = \begin{cases} 1 + c_{1,1}x_1 + c_{1,2}2x_1^2x_2^2 \\ 1 + c_{2,1}x_1t + c_{2,2}x_2t + c_{2,3}x_1x_2^2t^3. \end{cases} \quad (5.14)$$

By closer inspection, however, it is still not possible to identify the starting points because $\hat{q}_2(\mathbf{x}, t = 0) = 1$. This problem can be resolved according to [HS95, Lemma 3.1]: i.e., initial values are obtained by solving a binomial system for every cell that contributes to the mixed volume computation (i.e., for every gray cell in Figure 5.1). One can then increase t and obtain the solutions of the start system by tracking the solution paths to $\hat{\mathbf{Q}}(\mathbf{x}, t = 1) = \mathbf{Q}(\mathbf{x})$.

(iii) *Target System:* Finally we have all 4 solutions to $\mathbf{Q}(\mathbf{x})$. Now what remains is to construct a *linear* homotopy according to (4.21) and to increase t , starting at $t = 0$. At $t = 1$ the homotopy reduces to $\mathbf{F}(\mathbf{x})$ and provides the desired solutions of the target system.

5.4 Stability of Fixed Points

We proceed according to the usual procedure in dynamical systems (cf. Chapter 4) and, after computing the set of all fixed points, assess their (local) stability. A fixed point is locally stable if a neighborhood exists such that messages inside this neighborhood (i.e., messages sufficiently close to the fixed point), converge to the fixed point under the considered map [Tes12, pp.170].

Note that the Bethe free energy does not fully characterize the stability of a given fixed point (cf. Section 3.3.2): although stable fixed points are local minima of $\mathcal{F}_{\mathcal{B}}$, local minima must not be stable [Hes03]. A general way of investigating the stability of fixed points is available by the method of Lyapunov [Sch00, pp.93], though it is sufficient for all graphs considered in this work to restrict our analysis to linearization (the indirect method of Lyapunov) as introduced in Section 4.4.

This approach is well-established in the dynamical systems literature and it may consequently seem rather surprising that such a stability analysis has not been considered so far for BP. The main difficulty, however, is not the stability analysis as such, but the prerequisite of estimating all fixed points. Indeed, a local stability analysis has been performed for Ising models with vanishing local potentials [MK05], for which all fixed points are known. We can, however, rely on the NPHC method, as discussed in the preceding section, to solve the fixed point equations and to obtain the set of *all* fixed points. Subsequently, it is evident how to analyze the stability by computing the Jacobian and thus linearizing BP in every fixed point.

5.4.1 Reformulation of Belief Propagation for Binary Variables

Before we investigate the stability of all BP fixed points, we introduce an alternative parameterization of the update equations for the particular case of binary variables. This parameterization reduces the number of variables and eases some calculations without changing the properties of BP. Therefore, we express both messages along the same edge as a single message (cf. [MK05, KP17]) defined by

$$\nu_{ij}^n = \operatorname{atanh}(\mu_{ij}^n(X_j = 1) - \mu_{ij}^n(X_j = -1)). \quad (5.15)$$

Then, the update rule in (3.1) can be rewritten according to

$$\tanh(\nu_{ij}^{n+1}) = \tanh(J_{ij}) \tanh(h_{ij}^n), \quad (5.16)$$

where the cavity field h_{ij}^n acts on X_i , while neglecting the incoming message from X_j , according to

$$h_{ij}^n = \theta_i + \sum_{X_k \in \{\partial(i) \setminus X_j\}} \nu_{ki}^n. \quad (5.17)$$

Considering (5.16) and (5.17) (cf. [OW01]) it becomes evident that, for binary pairwise models, BP corresponds to the so-called cavity method [MPV87]. Its name stems from the fact that we essentially dig a cavity into the model by removing X_j . We then express the remaining field that acts X_i by the cavity field h_{ij} .

The marginals, or the mean, can then – similar as for BP – be computed by all incoming

messages according to

$$\begin{aligned} m_i &= \tanh(h_{ij} + \nu_{ji}) \\ &= \theta_i + \sum_{X_k \in \partial(i)} \nu_{ki}. \end{aligned} \quad (5.18)$$

In accordance with the consideration of BP as a dynamical system, we will denote the set of all re-parameterized messages by $\boldsymbol{\nu}$ and denote the mapping induced by BP as $\boldsymbol{\nu}^{n+1} = \mathcal{BP}(\boldsymbol{\nu}^n)$. Keeping the notation consistent, we say that BP converged to a fixed point

$$\boldsymbol{\nu}^\circ = \mathcal{BP}(\boldsymbol{\nu}^\circ), \quad (5.19)$$

if successive messages remain unchanged under BP. If, however, BP fails to converge, one can try to achieve convergence by one of the many modifications discussed in Section 3.5.1. One of these modifications that lends itself handsomely for our stability analysis is BP with damping.

Let us briefly state the update equations of BP with damping in terms of the reformulated messages again: $\mathcal{BP}_D(\boldsymbol{\nu}^n) = (1 - \epsilon)\mathcal{BP}(\boldsymbol{\nu}^n) + \epsilon\boldsymbol{\nu}^n$, where $\epsilon \in [0, 1)$ is the damping factor. Remember that, although the fixed points of BP with damping are fixed points of BP without damping as well (cf. (3.51)), the stability of the fixed points may change nonetheless.

5.4.2 Linearization

In order to assess the stability of the fixed points $\boldsymbol{\nu}^\circ$ we approximate $\mathcal{BP}(\cdot)$ by a linear function in its fixed point(s) and analyze the behavior of the linearized system. This is done by taking the partial derivatives of all messages, i.e., by analyzing the Jacobian matrix $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ with its elements defined as

$$\mathcal{F}'(\boldsymbol{\nu}^\circ)_{mn} = \frac{\partial \nu_{ij}^\circ}{\partial \nu_{kl}^\circ}, \quad (5.20)$$

where – given some ordering – (i, j) and (k, l) are the m^{th} and the n^{th} edge.

For binary pairwise models parameterized as in Section 5.4.1, the Jacobian is given as follows: without loss of generality²⁹ we consider only messages where $(i, j) \in \mathbf{E}$ and $(k, l) \in \mathbf{E}$ such that

$$\mathcal{F}'(\boldsymbol{\nu}^\circ)_{mn} = \begin{cases} \frac{\tanh(J_{ij})(1 - \tanh^2(h_{ij}))}{1 - \tanh^2(J_{ij}) \tanh^2(h_{ij})} & \text{if } i=l \text{ and } k \in \{\partial(i) \setminus X_j\} \\ 0 & \text{else.} \end{cases} \quad (5.21)$$

Let us briefly recap the properties of the eigenvalues with respect to the stability of the BP fixed points from Section 4.4. A fixed point $\boldsymbol{\nu}^\circ$ is locally stable if all eigenvalues have absolute value strictly smaller than one, i.e., if $\rho(\mathcal{F}'(\boldsymbol{\nu}^\circ)) < 1$ and BP converges if initialized sufficiently close enough. A fixed point is unstable with respect to BP if at least one eigenvalue exists outside the unit circle such that $\rho(\mathcal{F}'(\boldsymbol{\nu}^\circ)) > 1$. For $\rho(\mathcal{F}'(\boldsymbol{\nu}^\circ)) = 1$ stability of the nonlinear system cannot be inferred by just looking at the linear system.

In addition to damping, there are many other variants of BP available that have the same set of solutions $V_{\mathbb{R}_+}^*$ [WJW03a]; in particular, this includes different scheduling methods. Analyzing the stability of the fixed points under these variants, however, becomes problematic – mainly because the update function changes with time. Moreover, the fact that many messages are not

²⁹ If all combinations of m and n are considered we would effectively consider all $\frac{n(n-1)}{2}$ possible edges. All rows and columns that correspond to a message ν_{vw} , where $(v, w) \notin \mathbf{E}$ include only zero-values, however, and can therefore be neglected without changing the eigenvalues.

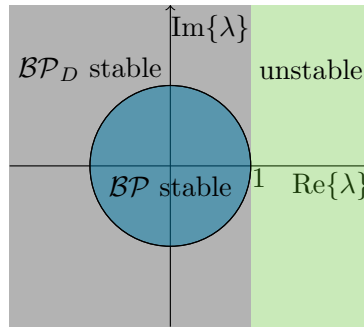


Figure 5.2: Eigenvalue spectrum of the Jacobian $\Lambda(\mathcal{F}'(\nu^\circ))$. BP is stable if all eigenvalues lie inside the blue region; BP with damping is stable if all eigenvalues lie inside the gray region; BP with and without damping is unstable if eigenvalues lie inside the green region.

updated, introduces eigenvalues with $\lambda_i = 1$, which renders the stability analysis by linearization impossible.

Therefore, of all the modified versions of BP, we will restrict our attention to damping. The application of damping modifies the eigenvalue spectrum according to

$$\Lambda(\mathcal{F}'_D(\nu^\circ)) = \Lambda(\mathcal{F}'(\nu^\circ)) \cdot (1 - \epsilon) + \epsilon. \quad (5.22)$$

Note that all eigenvalues are reduced by a factor $(1 - \epsilon)$ and experience a shift by ϵ along the real axis, i.e., $\text{Re}\{\Lambda(\mathcal{F}'_D(\nu^\circ))\} = \epsilon + \text{Re}\{(1 - \epsilon)\Lambda(\mathcal{F}'(\nu^\circ))\}$. A fixed point ν° is thus locally stable under $\mathcal{BP}_D(\cdot)$ if

$$\text{Re}\{\Lambda(\mathcal{F}'(\nu^\circ))\} < 1. \quad (5.23)$$

The correspondence between stability and the eigenvalue spectrum is summarized and visualized in Figure 5.2. A fixed point is stable under BP if all eigenvalues lie inside the unit circle (depicted by the blue area), under BP with damping if all eigenvalues have a real part strictly smaller than one (depicted by the gray area), and are unstable else (depicted by the green area).

Note that the properties of the Bethe Hessian (and thus the stability of the fixed points) further relate to some concepts from graph-theory; in particular to the Ihara zeta function [WF09] and – for models with vanishing local fields, i.e., where $\theta = 0$ – to the non-backtracking matrix (known as the Hashimoto-matrix) [SKZ14, SKZ17]. Note that the latter connection is insightful by any means but is only valid in a well-behaved region where all eigenvalues of the non-backtracking matrix are inside the unit circle (cf. [SKZ17, Theorem 2.1]).

5.5 Selected Models with Ising Potentials

Finally, we will now consider a range of probabilistic graphical models, apply the NPHC method to the fixed point equations (5.6), obtain all fixed points by first finding all isolated non-zero complex solutions, and analyze the stability of all fixed points.

We first evaluate and compare the accuracy of all fixed points obtained by NPHC; details for our evaluation criteria are presented in Section 5.6.1. Note that we will already anticipate some results from the later stability analysis and separately evaluate the accuracy of stable and unstable fixed points. However, here we rather ask *if* a given fixed point is stable, instead of *why* it is stable. Furthermore, we present the evolution of the fixed points over the parameter space in Section 5.6.5 and present some implications on the accuracy to better understand for which parameters BP can be expected to provide good results.

The capability of NPHC to obtain all fixed points, subsequently allows for a thorough stability




Figure 5.3: Considered Ising graphs: (a) grid-graph; (b) complete graph; (c) grid-graph with periodic boundary conditions.

analysis in Section 5.7. Our empirical observations further inspire some theoretical investigations on Ising models with unitary parameters in Section 5.8. There we show why convergence properties degrade with growing graph size and why strong local potentials help to achieve convergence. Finally, the performance of BP-decoding for error-correcting codes is analyzed in terms of the solution space in Section 5.9.

Considered Graphs

We consider different realizations of the Ising model with attractive, repulsive, and mixed interactions on a range of graphs. These graphs include complete graphs, grid-graphs, and grid-graphs with periodic boundary conditions (see Figure 5.3).

For the complete graph each pair of nodes is connected by an edge; it follows by definition that this is a regular graph, i.e., all variables $X_i \in \mathbf{X}$ have equal degree $d_i = N - 1$. Because of this, we can construct a Cayley tree to determine the phase-transitions.

The grid-graph has all edges aligned along the two-dimensional square lattice and, for finite-size graphs, contains variables of varying degrees. We additionally consider grid-graphs with periodic boundary conditions, where nodes on the boundary are joined by edges so that all variables have equal degree again.

5.6 Number of Fixed Points and Marginal Accuracy

This section provides an exhaustive solution space analysis of BP in terms of analyzing the properties for all fixed points. In order to obtain the set of all fixed points we have to solve the system of fixed point equations (5.6) first. The systems considered in this work are simply too large to be solved with symbolic methods or even with the NPHC method based on a linear homotopy. The BKK bound, however, takes into account the sparsity of the system $\mathbf{F}(\boldsymbol{\mu}, \boldsymbol{\alpha})$, induced by the graph structure, and reduces the number of solution paths to be tracked so that the problem can be solved in practice. We present a detailed runtime analysis in Section 5.6.6. Note that both the structure of (5.6) and the number of complex solutions in $V(\mathbf{F})$ remain the same if the graph structure is kept constant [CM17, CMN16]; depending on the potentials, however, the number of solutions in $V_{\mathbb{R}_+}^*(\mathbf{F})$ may change.

5.6.1 Evaluation Criteria

We evaluate all fixed points in terms of their marginal accuracy and compare them with marginals obtained by an implementation of BP without damping. Evaluation of the marginal

accuracy requires the availability of the exact marginals; these are obtained by the junction tree algorithm that we can still resort to because of the limited size of the considered models.

In particular, we evaluate the correctness of the approximated marginals in terms of their accuracy (cf. (3.38)). The expected mean (cf. (3.39)) provides an alternative way of evaluating the marginals and is particularly well suited for illustrative purposes. We will therefore not only compare different fixed points in terms of the MSE but also visualize the averaged mean of the exact solution $\langle m \rangle$ as well as of all BP fixed points $\langle \tilde{m} \rangle$. Remember that $\langle m \rangle - \langle \tilde{m} \rangle$ equals the average of all marginal errors (cf. (3.39)) and is thus well-suited for visually comparing the marginal accuracy of different fixed points.

We will not only compare different fixed points but also evaluate the effect of altering the evaluation function. A combination of properly weighted marginals may improve the marginal accuracy considerable but is often problematic in practice for the lack of methods that obtain more than a single fixed point. The NPHC method, however, computes the set of all solutions $\mathbf{V}_{\mathbb{R}_+^*}(\mathbf{F})$ and – by Theorem 1 – yields the set of all fixed points \mathbf{T} . This is the foundation for evaluating the marginal accuracy of different combinations.

To combine the marginals, we first need to compute the pseudomarginals \tilde{P}_B and the Bethe partition function \mathcal{Z}_B for all fixed points. Then, we utilize the obtained values of the Bethe partition function \mathcal{Z}_B together with the pseudomarginals \tilde{P}_B and modify the evaluation function \mathcal{D}_P in three different ways.

First, all fixed points are weighted by their partition function and are combined so that

$$\tilde{P}_B^T = \frac{1}{\sum_{\mathcal{Z}_B \in \mathbf{T}} \mathcal{Z}_B} \sum_{(\tilde{P}_B, \mathcal{Z}_B) \in \mathbf{T}} \tilde{P}_B \cdot \mathcal{Z}_B. \quad (5.24)$$

Second, we combine all fixed points that correspond to local minima of the Bethe partition function in the same way so that

$$\tilde{P}_B^M = \frac{1}{\sum_{\mathcal{Z}_B \in \mathbf{M}} \mathcal{Z}_B} \sum_{(\tilde{P}_B, \mathcal{Z}_B) \in \mathbf{M}} \tilde{P}_B \cdot \mathcal{Z}_B. \quad (5.25)$$

Note that these combinations are not just some heuristics. In particular the latter one reminds us of how the exact solution is expected to decompose according to the RSB assumption.

Finally, we aim to select one specific fixed point and hope that accurate marginals are to be found at the maximum of the Bethe partition function (i.e., the global minimum of the Bethe free energy). The fixed point maximizing \mathcal{Z}_B is consequently selected and evaluated as well; more formally it is given by

$$\tilde{P}_B^{MAX} = \arg \max_{\tilde{P}_B \in \mathbf{T}} \mathcal{Z}_B(\tilde{P}_B). \quad (5.26)$$

5.6.2 Grid Graphs with Random Factors (Spin Glasses)

Consider a grid graph of size $N = 3 \times 3$ with randomly distributed parameters. All pairwise and local potentials are sampled uniformly; i.e., $(J_{ij}, \theta_i) \sim \mathcal{U}(-K, K)$. The larger the support of the uniform distribution is, the more difficult the task of inference becomes; we choose $K = 3$, which is large enough to make BP fail to converge sometimes (cf. [SM07]).

According to (5.7) the system of equations consists of $M = 72$ equations in 72 unknowns. More specifically, (5.6) consists of 24 linear (i.e., normalization constraints), 40 quadratic, and 8 cubic equations; the total degree bounds the number of solutions by $d_t = 1^{24} \cdot 2^{40} \cdot 3^8 = 7.2 \cdot 10^{15}$. Tracking such an amount of solution paths is not feasible in practice, even with a parallel

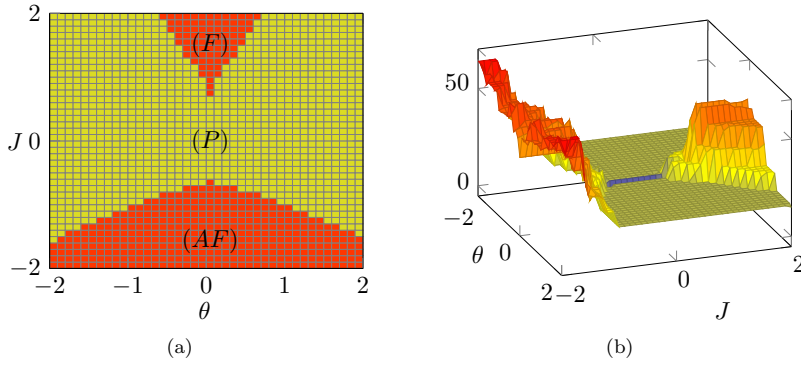


Figure 5.4: Number of fixed points on the grid graph of size $N = 3 \times 3$: (a) number of positive real fixed points (yellow: unique fixed point, red: three fixed points); (b) number of real fixed points. The increase in the number of both real solutions and positive real solutions, indicates a phase transition.

implementation of the NPHC method.

The system of equations in (5.6), however, is sparse. We can exploit this sparsity that is induced by the graph structure if we consider the BKK bound and reduce the computational complexity. The number of complex solutions for this graph is bounded by $\text{BKK} = 608$. After creating a suitable start system the problem is straightforward to solve with the NPHC method. It actually turns out that the BKK bound is tight for all graphs considered.

In particular we evaluate 100 grid graphs with random factors: on 99 graphs BP converged after at most 10^4 iterations. Although the grid graph has multiple loops and the constrained $\mathcal{F}_{\mathcal{B}}$ is not necessarily convex [Hes04, Corr.2], we observe that for all 100 graphs NPHC obtains a unique positive real solution that corresponds to a unique BP fixed point.

5.6.3 Grid Graphs with Uniform Factors

We further analyze BP on grid graphs of size $N = 3 \times 3$ with constant potentials among all nodes and edges; i.e., we specify the couplings for all edges by $J_{ij} = J$ and specify the fields for all nodes by $\theta_i = \theta$. We apply BP and NPHC for 1681 graphs in the parameter region $(J, \theta) \in \{-2, -1.9, \dots, 1.9, 2\}$ and illustrate the size of the solution set in Figure 5.4.

The number of solutions in $V_{\mathbb{R}_+}^*(\mathbf{F})$ is presented in Figure 5.4 (a). In the well-behaved region (P) of the parameter space a unique fixed point exists, whereas three fixed points exist in (F) and (AF) – this is in accordance with statistical mechanics [MM09, p.43].³⁰ Interestingly, we observe a close relation between the onset of phase transitions and the increase in the number of real solutions in Figure 5.4 (b). Most of these solutions, however, correspond to negative message values that violate Lemma 1 and are not feasible.

BP converges to some fixed point on all 1681 graphs within at most 10^4 iterations. This raises a couple of questions: What is the approximation error of BP if it converges to the most accurate fixed point? Or, speaking in terms of free energies, how large is the gap between the global minimum of the constrained $\mathcal{F}_{\mathcal{B}}$ and the minimum of the Gibbs free energy? To answer this question we evaluate the correctness of the approximated marginals by computing the MSE between the exact and the approximated marginals according to (3.38). The results are presented in Table 5.1. Averaged over all graphs we can see that BP does not necessarily converge to the most accurate fixed point. For the NPHC method we present the MSE for the fixed point with the lowest MSE; this highlights the existence of fixed points, which give *more* accurate approximations than BP. Looking at all parameter regions separately we can see that BP does converge to the global optimum in (P), as well as in (F). In the antiferromagnetic

³⁰ Note that the graph under consideration is of finite size and thus d_i varies among the nodes. As a consequence the partitioning according to (5.3) - (5.5), is only an approximation.

Table 5.1: MEAN SQUARED ERROR, I.E., $E_P(m)$ OF MARGINALS AND COMBINED MARGINALS OBTAINED BY BP AND NPHC FOR THE GRID GRAPH WITH UNITARY FACTORS.

Parameters		Fixed Points		Combined		
Couplings	Local Field	BP	NPHC	MAX	ALL	STABLE
$J \in [-2, 2]$	$\theta \in [-2, 2]$	0.197	0.016	0.037	0.005	0.004
$J \in (F)$	$\theta \in (F)$	0.010	0.010	0.076	$9.2 \cdot 10^{-4}$	$1.0 \cdot 10^{-9}$
$J \in (AF)$	$\theta \in (AF)$	0.836	0.050	0.126	0.003	$1.5 \cdot 10^{-6}$
$J \in (P)$	$\theta \in (P)$	0.006	0.006	0.006	0.006	0.006

region (AF) , BP converges to a fixed point that does not necessarily give the best possible approximation.

If we consider regions with multiple fixed point solutions (i.e., (F) and (AF)) it becomes obvious that the fixed point maximizing the partition function (i.e., \tilde{P}_B^{MAX}) is not necessarily the best one; this is especially surprising as BP obtains the best possible fixed point solution inside (F) .

For $\theta = 0$ it turns out that initializing all messages to the same value $\mu_{ij}^1(x_j) = \mu^1 = \frac{1}{|\mathcal{X}|}$ will result in a fixed point where $\tilde{P}_{X_i}(x_i) = 0.5$ (and $m_i = 0$) for all $X_i \in \mathbf{X}$. Although this fixed point is identical to the exact one, it may be unstable (cf. Figure 5.6).

Inspired by these observations, one should not only consider the fixed point maximizing \mathcal{Z}_B (i.e., \tilde{P}_B^{MAX}), but rather obtain multiple fixed points by NPHC and combine them. Indeed, especially inside region (F) a combination of *all* fixed point solutions according to (3.53), i.e., \tilde{P}_B^T increases the accuracy of the approximation. If we combine the fixed points at local minima only (which are all stable in this specific case), i.e., \tilde{P}_B^M , the accuracy increases even more and gives the *most* accurate approximation over the entire parameter space.

5.6.4 Complete Graphs with Uniform Factors

We consider a complete graph with $N = 4$ binary random variables and illustrate the number of solutions in Figure 5.5. The system of equations (5.6) consists of 36 equations in 36 unknowns and has its number of solutions bounded by the total degree $d_t = 1^{12} \cdot 2^{24} = 16.8 \cdot 10^6$. Similar as for the grid graph, a much tighter bound of $\text{BKK} = 120$ is provided by the BKK bound. Among all four nodes, we apply unitary factors, specified by $J_{ij} = J$ and $\theta_i = \theta$. This type of graph is particularly interesting because one can derive exact conditions where phase transitions occur (cf. Section 5.2.1).

For $(J, \theta) \in (P)$, BP has a unique fixed point, which is a stable attractor in the whole message space [MK05]. In (F) three fixed points satisfy (5.6), one of which is unstable and a local minimum of \mathcal{Z}_B . Both other fixed points are local maxima of \mathcal{Z}_B and BP converges to one of them.

For repulsive models, BP only converges inside (P) and not inside (AF) as shown in Figure 5.5 (a). We can see two interesting effects: first, in Figure 5.5 (c) the number of real solutions increases at the onset of phase transitions; second, even though the convergence of BP breaks down at the phase transition, a unique fixed point exists inside (AF) (Figure 5.5 (b)) that gives an accurate approximation (cf. Table 5.2).

Similar as in Section 5.6.3, we assess the MSE of the marginals obtained by NPHC and BP; the results – averaged over all graphs, and for each distinct region – are presented in Table 5.2. Indeed, inside region (AF) the marginals obtained by NPHC give a much better approximation than BP does.




Figure 5.5: Fully connected graph with $N = 2 \times 2$: (a) convergence of BP: for the blue region BP did convergence – for the red region it did not converge after $4 \cdot 10^5$ iterations; (b) number of fixed points (yellow: unique fixed point, red: three fixed points); (c) number of real solutions – note the sudden increase at the onset of phase transitions.

Table 5.2: MEAN SQUARED ERROR, I.E., $E_P(m)$ OF MARGINALS AND COMBINED MARGINALS OBTAINED BY BP AND NPHC FOR THE FULLY CONNECTED GRAPH WITH UNITARY FACTORS.

Parameters		Fixed Points		Combined		
Couplings	Local Field	BP	NPHC	MAX	ALL	STABLE
$J \in [-2, 2]$	$\theta \in [-2, 2]$	0.069	0.007	0.011	0.003	0.0027
$J \in (F)$	$\theta \in (F)$	0.034	0.033	0.070	0.002	$2.0 \cdot 10^{-8}$
$J \in (AF)$	$\theta \in (AF)$	0.304	0.004	0.004	0.004	0.004
$J \in (P)$	$\theta \in (P)$	0.003	0.003	0.003	0.003	0.003

Furthermore, note that the fixed point maximizing \mathcal{Z}_B does not always give the best approximation. For $\theta = 0$, if all messages are initialized to the same value, BP obtains the exact marginals (cf. Section 5.6.3). If multiple fixed points exist, a weighted combination of all marginals according to (3.53) increases the accuracy – only considering stable solutions, i.e., \tilde{P}_B^M , gives the *most* accurate approximations.

5.6.5 Fixed Point Evolution

To better understand the influence of the parameters (J, θ) we specifically investigate how they affect the accuracy of the fixed points. Therefore, we fix the values of $\theta \in \{0, 0.1, 0.5\}$, vary $J \in [-2, 2]$, and compare the fixed point solutions obtained by NPHC to the exact solution. For illustrative purposes we consider the mean, averaged over all variables, of the exact solution $\langle m \rangle$ and of the approximate solution $\langle \tilde{m} \rangle$ for all fixed points obtained by NPHC; the results are illustrated for both the grid graph and the complete graph in Figure 5.6 and Figure 5.7.

The exact solution (red) is obtained by the junction tree algorithm [LS88]; solutions to (5.6) are obtained by NPHC and are depicted by blue dots (stable) and by green or black dots (unstable). We further emphasize the fixed point that maximizes \mathcal{Z}_B , i.e., \tilde{P}_B^{MAX} , in orange.

We will already discuss some implications of the local stability analysis here, but focus mainly on the accuracy of the fixed points and defer a thorough discussion of the stability analysis to Section 5.7.

All fixed points for a grid graph with $N = 3 \times 3$ binary random variables are shown in

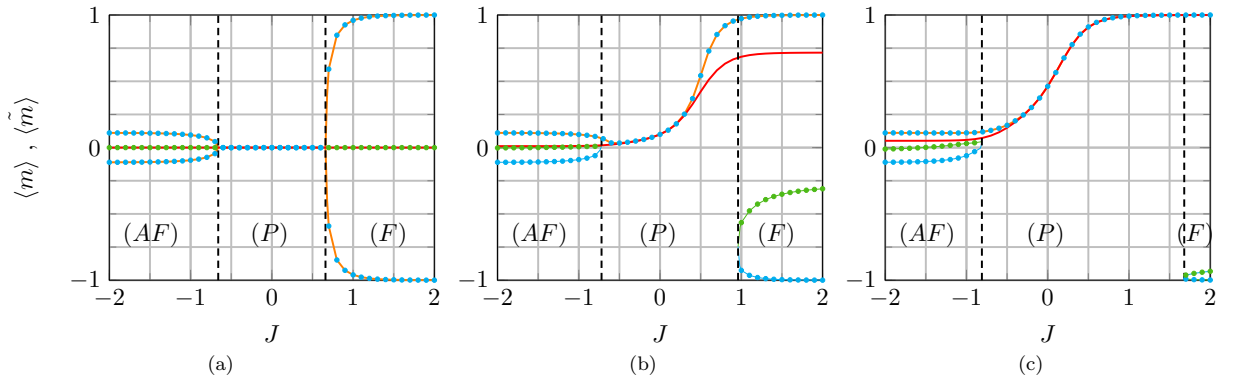


Figure 5.6: Results for the grid graph of size $N = 3 \times 3$; average mean $\langle m \rangle$ and $\langle \tilde{m} \rangle$ for $J \in [-2, 2]$ and for: (a) $\theta = 0$, (b) $\theta = 0.1$, and (c) $\theta = 0.5$. The exact solution is illustrated in red. All fixed points obtained by NPHC are depicted by blue dots (stable) and by green dots (unstable). The fixed point maximizing the partition function is illustrated in orange.

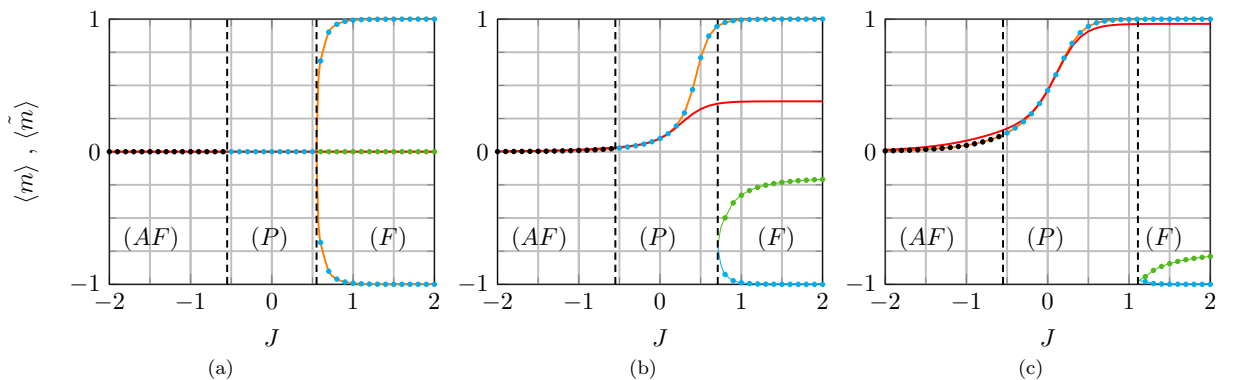


Figure 5.7: Results for the complete graph of size $N = 2 \times 2$; average mean $\langle m \rangle$ and $\langle \tilde{m} \rangle$ for $J \in [-2, 2]$ and for: (a) $\theta = 0$, (b) $\theta = 0.1$, and (c) $\theta = 0.5$. The exact solution is illustrated in red. All fixed points obtained by NPHC are depicted by blue dots (stable) and by green/black dots (unstable). The fixed point maximizing the partition function is illustrated in orange.

Figure 5.6. The worst performance in terms of accuracy is observed for $\theta = 0$. Despite the existence of one fixed point that corresponds to the exact solution with $\langle m \rangle = \langle \tilde{m} \rangle = 0$, BP fails to provide accurate marginals as this particular fixed point is only stable inside (P). As the coupling strength $|J|$ increases to the onset of phase transitions, two additional fixed points emerge. These additional solutions are symmetric, stable, and guarantee the convergence of BP on this graph (see Figure 5.6 (a)).

For $\theta \neq 0$ a unique stable fixed point exists inside (P). If we gradually increase J until $(J, \theta) \in (F)$ two additional fixed points emerge, one of which is unstable (see Figure 5.6 (b)). Note that the fixed point maximizing \mathcal{Z}_B (orange) remains stable for all values of $J \in [-2, 2]$. An increase in θ (see Figure 5.6 (c)) does enlarge the region where a unique fixed point exists and further increases the accuracy of the fixed point maximizing \mathcal{Z}_B . For $(J, \theta) \in (AF)$ a similar behavior is observed; i.e., for small values of θ the unstable fixed point has the highest accuracy, but as θ increases, the accuracy of the fixed point maximizing \mathcal{Z}_B increases as well.

All fixed points for the complete graph with $N = 4$ binary random variables are shown in Figure 5.7. For $\theta = 0$ and large values of J the fixed point with $\langle m \rangle = \langle \tilde{m} \rangle = 0$ is unstable and is accompanied by two symmetric, stable fixed points (Figure 5.7 (a)). In contrast to the grid graph a unique and accurate fixed point exists for $(J, \theta) \in (AF)$; yet, this is of no avail as the fixed point is unstable (see Figure 5.5 (a)). This highlights that the existence of a unique, accurate, fixed point does not necessarily imply convergence of BP.

Table 5.3: RUNTIME COMPARISON BETWEEN BP AND NPHC (IN SECONDS).

	Grid Graph: Random (J, θ)		Grid Graph: Unitary (J, θ)		Fully Conn. Graph: Unitary (J, θ)	
	total	parallel	total	parallel	total	parallel
Mixed Vol.	1364.5	11.2	1311.1	12.0	0.27	–
Path Track.	69.0	0.97	70.5	1.0	3.77	–
Post Proc.	29.9	1.40	43.6	4.1	2.3	–
NPHC	13.57		17.1		6.34	
BP	$3.6 \cdot 10^{-3}$		$0.7 \cdot 10^{-3}$		0.03	

For $\theta \neq 0$ the non-convergent region (AF) is slightly reduced, but the problem of a unique unstable fixed point persists. In contrast to the grid graph, the complete graph contains odd-length cycles, thus allowing for frustrations if all edges are repulsive. This points at a close connection between the existence of frustrations and the existence of a unique unstable solution. The complete graph behaves similar to the grid graph for $(J, \theta) \in (F)$, i.e., the accuracy of the fixed point maximizing \mathcal{Z}_B increases as θ increases (see Figure 5.7 (b) and Figure 5.7 (c)).

Our main findings are: First, increasing the field θ increases the accuracy of the fixed point maximizing \mathcal{Z}_B . Second, for $\theta \neq 0$ the fixed point maximizing \mathcal{Z}_B is unique and varies continuously under a change of J . Finally, the two stable fixed points are close to being symmetric, i.e., $\tilde{P}_{X_i}^{(1)}(x_i) \cong 1 - \tilde{P}_{X_i}^{(2)}(x_i)$. Consequently combining both stable fixed points will not lead to good approximations unless a proper weighting by \mathcal{Z}_B is applied. Applying a proper weighting, however, leads to accurate approximations (cf. Table 5.1- 5.2).

5.6.6 Runtime Analysis

The time required for solving (5.6) is presented in Table 5.3 for grid graphs with random factors (Section 5.6.2), grid graphs with unitary factors (Section 5.6.3), and fully connected graph with unitary factors (Section 5.6.4). Comparing the overall computation time of the NPHC method to BP it becomes obvious that NPHC is no alternative in terms of computational efficiency. It is, however, the only method that is guaranteed to obtain *all* fixed points – we were not able to apply the Gröbner basis method beyond a single-cycle graph with $N = 4$. For our computations we utilized a cluster-system with 160 CPUs.

If we compare the overall computation time to the actual computation time utilizing the parallel implementation it becomes obvious that NPHC benefits tremendously from the high degree of parallelization. The runtime of BP depends mainly on the number of iterations and less on the size of the graph. Consequently, the stability of fixed points directly affects the performance of BP (cf. non-convergent region in Figure 5.5 (a)). The NPHC method is much less sensitive to the stability of fixed point solutions; the mixed volume computation, which has the largest influence on the overall runtime, rather depends on the number of variables in (5.6). Note the mixed volume computation does not depend on the parameters. If one is interested in the fixed points for different parameter-sets on the same graph it would suffice to compute the mixed volume and the start system only once; we did not do this to allow for a fair comparison.

5.7 Empirical Stability Analysis of Belief Propagation on Ising Models

So far we have been rather sloppy with the notion of stability and assumed knowledge about the stability of the fixed points without discussing how to gain this knowledge. The fact that we have obtained all fixed points, irrespective of their properties under a specific map, admits a thorough stability-analysis for all of them.

In this section, we analyze the local stability of all fixed points for the models discussed so far. Then, we compare the results to known results of both infinite and finite-size graphs with vanishing local fields $\theta = 0$. We restrict our analysis to attractive and repulsive models (cf. Section 2.4.3), because this allows us to change the behavior of BP with just a single parameter.

We briefly present the results of [MK05] for vanishing fields in Section 5.7.1. Then we extend the analysis to graphs with non-vanishing fields, discuss some empirical observations, and interpret the implications. Note that the results of our stability analysis are also illustrated in terms of the average mean $\langle \tilde{m} \rangle$ over the couplings-strength in Figure 5.6 (for the 3×3 grid graph) and 5.7 (for the complete graph of size $N = 4$). All fixed points of BP are colored according to the discussion of the eigenvalue spectrum (cf. Figure 5.2): fixed points are depicted in blue if stable under $\mathcal{BP}(\cdot)$, in black if stable under $\mathcal{BP}_D(\cdot)$, and in green if unstable. For reference we also illustrate the exact solution in red. A more formal analysis is presented in Section 5.8.

5.7.1 Vanishing Local Field

It is generally assumed that the case of vanishing local fields is the worst-case scenario [MK05]; we will confirm this assumption empirically and analytically. For $\theta = 0$ a trivial exact solution exists – namely marginals that are uniform over all states, i.e., for all $X_i \in \mathbf{X}$ the singleton marginals are $\tilde{P}_{X_i}(x_i) = 0.5$ and the mean is $m_i = 0$.

For attractive models with $J > 0$ sufficiently small, BP converges to this trivial (paramagnetic) fixed point, which is unique and stable. As the coupling-strength increases, the eigenvalue with the largest magnitude λ_{max} increases as well (see Theorem 4) and as λ_{max} crosses the unit circle the paramagnetic fixed point remains unchanged but becomes unstable. At the same time, two additional fixed points appear – these fixed points are symmetric and stable (cf. Figure 5.6 (a) and Figure 5.7 (a)).

For repulsive models with $J < 0$ all entries of the Jacobian swap in sign, i.e., $\mathcal{F}'_-(\nu^\circ) = -\mathcal{F}'_+(\nu^\circ)$, where $-$ and $+$ indicate the Jacobian for repulsive and attractive models, respectively. It follows that $\Lambda(\mathcal{F}'_-(\nu^\circ)) = -\Lambda(\mathcal{F}'_+(\nu^\circ))$. Consequently, the local stability of the fixed point is invariant under a sign-change of J and instability occurs precisely for the same coupling strength as before (cf. Figure 5.6 (a) and Figure 5.7 (a)). There is one important difference though – the dominant eigenvalue is negative now, such that damping helps to achieve convergence.

5.7.2 Non-Vanishing Local Field

For reasons of simplicity we restrict our analysis to models with a single value $\theta_i = \theta \neq 0$ for all variables. Because the Ising model is symmetric with respect to the local fields θ_i , it is sufficient to consider only non-negative local fields. Moreover, the same qualitative results hold if we allow for $\theta_i \geq 0$ in general.

Attractive Models

For small values of $J > 0$ a unique fixed point exists to which BP converges and for which all eigenvalues lie inside the unit circle. If we gradually increase the coupling-strength to the point where instability occurred for $\theta = 0$, we observe that all eigenvalues are still inside the




Figure 5.8: Eigenvalue spectra of the unstable fixed point of the complete graph with $N = 4$. If all eigenvalues are inside the unit-circle BP converges without damping. If no eigenvalues lie on the right-hand side of the vertical line BP converges with damping. (a) $J = 1.5$ and $\theta = 0$: notice that $\text{Re}\{\lambda_{max}\} > 1$, i.e., damping does not help; (b) $J = 1.5$ and $\theta = 0.5$: notice how the external field reduces the magnitude of the eigenvalues (cf. Theorem 3); (c) $J = -1.5$ and $\theta = 0.5$: the fixed point is unstable but can be stabilized with damping because $\text{Re}\{\lambda_i\} < 1$.




Figure 5.9: Eigenvalue spectra of the grid-graph with $N = 9$; all eigenvalues are symmetric because the graph is bipartite (cf. Theorem 5). (a) $J = 1.5$, and $\theta = 0$: the fixed point is unstable and damping does not help; (b) $J = 1.5$ and $\theta = 0.5$: only a unique stable fixed point exists (cf. Figure 5.6 (c)); note the qualitative difference in the spectrum; (c) $J = -1.5$ and $\theta = 0.5$: because of the symmetric spectrum, damping does not help in for repulsive models.




Figure 5.10: Eigenvalue spectra for $J = 1.5$ and $\theta = 0.5$ of the (a) grid-graph with $N = 9$ and periodic boundary conditions; (b) grid-graph with $N = 16$.

unit circle, i.e., the spectral radius $\rho(\mathcal{F}'(\boldsymbol{\nu}^\circ)) < 1$; the fixed point is consequently stable (cf. Theorem 3 for more details). If we further increase θ , the spectral radius decreases and a unique stable fixed point persists for even larger values of J (compare Figure 5.6 (b) and Figure 5.7 (b) with Figure 5.6 (c) and Figure 5.7 (c)).

Now, if we increase J – although the fixed point remains stable – two additional fixed points emerge beyond some critical point $J_C(\mathcal{G}, \theta)$ (see Figure 5.6 (b) and 5.7 (b)). As opposed to vanishing local fields, the unique fixed point is continuously deformed and remains stable though. The second stable fixed point corresponds to some self-preserving state of magnetization³¹ and is accompanied by another unstable fixed point (Figure 5.8 (b)).

The graph structure is supposed to influence the spectral radius as well. This becomes obvious if we enlarge the size of a graph while keeping its local structure unchanged, e.g., by increasing the grid-graph from $N_1 = 9$ to $N_2 = 16$. A comparison of Figure 5.9 with Figure 5.10 reveals an increased spectral radius, i.e., $\rho(\mathcal{F}'_{N_1}(\boldsymbol{\nu}^\circ)) < \rho(\mathcal{F}'_{N_2}(\boldsymbol{\nu}^\circ))$ (cf. Theorem 4). It is interesting that the largest eigenvalue – and its symmetric counterpart for bipartite graphs – are the only eigenvalues that experience a relevant increase in their real part. The real part of all other eigenvalues $\lambda_i \in \Lambda(\mathcal{F}'(\boldsymbol{\nu}^\circ)) \setminus \{\lambda_{\max} : |\operatorname{Re}\{\lambda_{\max}\}| = \rho(\mathcal{F}'(\boldsymbol{\nu}^\circ))\}$ is bounded by $|\operatorname{Re}\{\lambda_i\}| < 1$.

In terms of accuracy, BP performs better for models with $\theta \neq 0$; compared to models with $\theta = 0$, a stable fixed point exists that lies closer to the exact solution. Loosely speaking, increasing the local fields effectively reduces the influence of the couplings. This does not only lead to better convergence properties of BP (cf. Corollary 4.1), but reduces the approximation-error as well.

Repulsive Models

For small values of $J < 0$ a unique fixed point exists for which all eigenvalues lie inside the unit circle. If we further decrease J , a change of behavior can be observed beyond some critical value $J_C(\mathcal{G}, -\theta)$. Interestingly, $|J_C(\mathcal{G}, -\theta)| \leq |J_C(\mathcal{G}, \theta)|$ with equality if and only if $\theta = 0$. In Figure 5.8 (b) and 5.8 (c) we can also see that $\rho(\mathcal{F}'_-(\boldsymbol{\nu}^\circ)) > \rho(\mathcal{F}'_+(\boldsymbol{\nu}^\circ))$, i.e., the invariance of local stability under sign-change of J does not hold in general.

Let $J < J_C(\mathcal{G}, -\theta)$, then it depends on the graph structure whether a unique fixed point exists and if damping is useful. It turns out that these two properties are closely connected:

First, consider the complete graph with $N = 4$ or the grid-graph with periodic boundary conditions; these models have a unique fixed point but frustrations exist because of cycles with odd-length. Consequently $\Lambda(\mathcal{F}'(\boldsymbol{\nu}^\circ))$ is non-symmetric (cf. Theorem 5). Besides the sign-change, the spectral radius increases as well so that $\rho(\mathcal{F}'_-(\boldsymbol{\nu}^\circ)) > \rho(\mathcal{F}'_+(\boldsymbol{\nu}^\circ))$; the dominant eigenvalue, however, has a negative sign and $\mathbb{R}(\Lambda(\mathcal{F}'(\boldsymbol{\nu}^\circ))) < 1$ so that an appropriate damping term exists that enforces BP to converge.

Second, consider both grid-graphs with $N_1 = 9$ and $N_2 = 16$; these graphs are bipartite and have a symmetric spectrum $\Lambda(\mathcal{F}'(\boldsymbol{\nu}^\circ))$ (see Figure 5.9, Figure 5.10 (b), and Theorem 5). Because of the symmetric spectrum BP behaves similarly as in the attractive case: i.e., no damping term exists that would stabilize the unstable fixed point; two additional fixed points exist, however, that are stable.

Multiple fixed points for repulsive models exist if and only if the underlying graph is bipartite. Such graphs can be decomposed into two disjoint subsets $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z}$ such that the means follow a "checkerboard-distribution" where for all $X_i \in \mathbf{Y}$ the mean is $m_i > 0$ and for all $X_j \in \mathbf{Z}$ the mean is $m_j < 0$, or exactly the other way round. This explains the existence of two stable fixed points. Note that the according average mean $\langle \tilde{m} \rangle$ behaves as follows: (i) if N is even both stable fixed points are symmetric and have the same average mean; (ii) if N is odd both stable fixed points have different average means, because the subsets differ in size, i.e., $|\mathbf{Y}| \neq |\mathbf{Z}|$. The

³¹ Self-preserving states are stable fixed points where $\langle \tilde{m} \rangle$ points into the opposite direction as θ .

difference in $\langle \tilde{m} \rangle$ reduces as the number of variables increases. In the limit of $N = \infty$, the average means of all fixed points collapse onto the same value.

To conclude our observations: either a unique fixed point exists, which may be unstable but can be stabilized by damping (black); or a fixed point exists which is unstable under any form of damping (green) but is accompanied by two stable fixed points (blue) (cf. Figure 5.6 and Figure 5.7). Therefore, non-vanishing fields increase the accuracy of BP (Figure 5.6 and Figure 5.7) and lead to better convergence properties (Theorem 3) in all experiments.

5.8 Theoretical Stability Analysis

Here we present some more formal arguments and explain the observations made in Section 5.7. We start with the Perron-Frobenius Theorem and specify some implications for the particular form of $\mathcal{F}'(\boldsymbol{\nu}^\circ)$. Then we provide properties of the eigenvalue spectrum $\Lambda(\mathcal{F}'(\boldsymbol{\nu}^\circ))$ and explain the influence of finite-size graphs in non-vanishing local fields. We consider only connected graphs (cf. Section 2.2), where all X_i have minimum degree $\min d_i \geq 2$. Note that it is straightforward to absorb any X_i with $d_i = 1$ into its neighbor.

Lemma 2. *The Jacobian matrix $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ of a connected graph G is irreducible if $\min d_i \geq 2$.*

Proof. The adjacency matrix of a connected undirected graph is irreducible. Let us construct the Jacobian $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ as in (5.20); note that we can partition $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ into $N \times N$ block matrices $[\mathcal{F}'(\boldsymbol{\nu}^\circ)]_{ik}$ of size $d_i \times d_k$ each. These blocks contain non-zero values if and only if $(i, j), (k, l) \in \mathbf{E}$ for $X_j, X_k \in \mathbf{X}$; i.e., if $a_{ik} = 1$. It follows that $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ is irreducible as well. \square

Theorem 2 (Perron Frobenius Theorem). *A real non-negative square matrix \mathbf{B} has its spectral radius bounded as follows:*

$$\min_i \sum_j b_{ij} \leq \rho(\mathbf{B}) \leq \max_i \sum_j b_{ij}. \quad (5.27)$$

If \mathbf{B} is irreducible the largest eigenvalue is a positive real number $\lambda_{max} = \rho(\mathbf{B})$.

Corollary 2.1 (Implications for regular graphs). *A regular graph has $d_i = d$ for all $X_i \in \mathbf{X}$. If the couplings and fields are constant, $\sum_n \mathcal{F}'(\boldsymbol{\nu}^\circ)_{mn} = c$ is constant for all rows. By the Perron Frobenius Theorem and by Lemma 2 it follows that $\lambda_{max} = c$.*

Consider an infinite-size grid graph with unitary couplings $J > 0$. Under these assumptions it is fairly straightforward to provide theoretical insights that explain why the existence of non-vanishing fields enhances the convergence properties. We generalize the results subsequently, as the same argument can be applied to finite-size graphs with different J_{ij} .

Theorem 3. *Let G_∞ be an infinite-size graph with $d_i = d$ and purely attractive interactions $J > 0$. Then, the existence of a non-vanishing external field $\theta \neq 0$ stabilizes BP.*

Proof. First assume a vanishing external field, i.e., $\theta = \theta_0 = 0$, in which case the trivial fixed point has identical messages $\mu_{ij}^\circ(x_j) = 0.5$ so that $\nu_{ij}^\circ = 0$. Consequently, (5.21) reduces to

$$\mathcal{F}'_{\theta_0}(\boldsymbol{\nu}^\circ)_{mn} = \begin{cases} \tanh(J) & \text{if } i = l \text{ and } k \in \{\partial(i) \setminus X_j\} \\ 0 & \text{else.} \end{cases} \quad (5.28)$$

Because all variables $X_i \in \mathbf{X}(G_\infty)$ have equal degree $\sum_n \mathcal{F}'(\boldsymbol{\nu}^\circ)_{mn} = c$ for all m . It follows from (5.21) and the Perron-Frobenius Theorem that $\rho(\mathcal{F}'_{\theta_0}(\boldsymbol{\nu}^\circ)) = \tanh(J) \cdot (d_i - 1)$. Without

loss of generality we exploit symmetry properties and assume an external field $\theta_\delta > 0$. By (3.1) and (5.16) it becomes obvious that every fixed point message $\nu_{ij}^\circ \neq 0$ and $h_{ij} \neq 0$. Note that qualitatively, it does not matter whether h_{ij} is positive or negative as we only consider $\tanh^2(h_{ij})$.

Because $0 \leq \tanh(J) < 1$ it follows that $\tanh(J) > \tanh^2(J)$ and consequently, as

$$\mathcal{F}'_{\theta_\delta}(\boldsymbol{\nu}^\circ)_{mn} = \begin{cases} \frac{\tanh(J)(1-\tanh^2(h_{ij}))}{1-\tanh^2(J)\tanh^2(h_{ij})} & \text{if } i = l \text{ and } k \in \{\partial(i) \setminus X_j\} \\ 0 & \text{else,} \end{cases} \quad (5.29)$$

all non-zero entries of $\mathcal{F}'_{\theta_0}(\boldsymbol{\nu}^\circ)_{mn}$ are element-wise larger than $\mathcal{F}'_{\theta_\delta}(\boldsymbol{\nu}^\circ)_{mn}$. Let us define a field-dependent scaling term $\kappa_\theta \in (0, 1)$ of the Jacobian matrix, then $\mathcal{F}'_{\theta_\delta}(\boldsymbol{\nu}^\circ) = \kappa_\theta \mathcal{F}'_{\theta_0}(\boldsymbol{\nu}^\circ)$.

Loosely speaking the existence of some non-vanishing field reduces all entries of the Jacobian matrix and consequently reduces the spectral radius. \square

Now choose a critical value J_C at the onset of instability³² so that $\rho(\mathcal{F}'_{\theta_0}(\boldsymbol{\nu}^\circ)) = \lim_{\epsilon \rightarrow 0} (1 + \epsilon)$. If the external field dampens the spectral radius, so that $\rho(\mathcal{F}'(\boldsymbol{\nu}^\circ)) < 1$ the unstable fixed point vanishes, and only a unique stable fixed point remains. Interestingly, this observation holds in all experiments, i.e., for attractive models an unstable fixed point exists if and only if multiple fixed points are present.

Suppose we either have purely attractive or purely repulsive interactions, then all entries of the Jacobian matrix have the same sign as the couplings and are bounded.

Lemma 3. *All entries of the Jacobian are bounded by $\mathcal{F}'(\boldsymbol{\nu}^\circ)_{mn} \in [0, 1]$ if $J_{ij} > 0$ and by $\mathcal{F}'(\boldsymbol{\nu}^\circ)_{mn} \in (-1, 0]$ if $J_{ij} < 0$.*

Proof. If all couplings have the same sign, the qualitative result of (5.29) holds as well. For purely repulsive interactions we just have to swap signs. \square

Theorem 4. *The finite-size manifestations of phase transitions occur beyond the theoretical phase transitions. Let us define a parameter-set (J_C, θ_C) for which λ_{max} crosses the unit-circle. Then, for a graph G_N with N nodes, the values of the critical parameters (J_C, θ_C) decrease as N increases.*

Consider two finite-size graphs with identical structure³³ with different size $N_1 < N_2$, then $\rho(\mathcal{F}'_{N_1}(\boldsymbol{\nu}^\circ)) \leq \rho(\mathcal{F}'_{N_2}(\boldsymbol{\nu}^\circ)) \leq \rho(\mathcal{F}'_\infty(\boldsymbol{\nu}^\circ))$.

Proof. Assume $J > 0$ and w.l.g. $\theta = 0$ (for the influence of non-vanishing external field see Theorem 3). Then on G_∞ – or for any other grid-graph with periodic boundary conditions – the degree $d_i = d$ is constant for all $X_i \in \mathbf{X}$. It follows from Corollary 2.1 that the largest eigenvalue is given by $\lambda_{max} = (d_i - 1) \cdot \tanh(J)$.³⁴ Suppose we increase the couplings to $J^{new} > J$ and denote the change of its parameters as $\kappa = \frac{\tanh(J^{new})}{\tanh(J)}$. It is obvious that $\lambda_{max}^{new} = \kappa \cdot \lambda_{max}$.

Suppose we have some finite-size graph where d_i is not constant but depends on X_i . If the couplings increase as before each row of $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ experiences a different amount of scaling. In all generality this is described by

$$\begin{bmatrix} c_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & c_K \end{bmatrix} \cdot \mathcal{F}'(\boldsymbol{\nu}^\circ) \quad (5.30)$$

³² If $\lambda_{max} = 1$ we cannot infer from the linearized to the nonlinear map [Tes12]. We introduce an ϵ -small term to avoid this subtlety.

³³ For example two-dimensional grid-graphs with $N = 3 \times 3$ or $N = 4 \times 4$ variables.

³⁴ For vanishing fields, i.e., if $\theta_i = 0$, the Jacobian matrix is fully defined by J_{ij} .

where c_k depends on d_i and $K = \sum_{i=1}^N d_i$. Let us reformulate (5.30) in all detail to

$$\mathbf{W} = \sum_{m=2}^{\max(d_i)} \mathbf{C}_m \mathcal{F}'(\boldsymbol{\nu}^\circ), \quad (5.31)$$

where \mathbf{C}_m is a diagonal matrix with values $\mathbf{C}_{m,kk} = \kappa_m$ if the k^{th} line of $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ corresponds to a variable with $d_k \geq m$ and 0 otherwise. Since the largest eigenvalue is a positive real number it follows that the largest eigenvalue is the sum of the individual eigenvalues, so that

$$\rho(\mathbf{W}) = \sum_{m=2}^{\max(d_i)} \rho(\mathbf{C}_m \mathcal{F}'(\boldsymbol{\nu}^\circ)). \quad (5.32)$$

We still have to show that

$$\rho(\mathcal{F}'_{N_1}(\boldsymbol{\nu}^\circ)) \stackrel{(a)}{\leq} \rho(\mathcal{F}'_{N_2}(\boldsymbol{\nu}^\circ)) \stackrel{(b)}{\leq} \rho(\mathcal{F}'_\infty(\boldsymbol{\nu}^\circ)), \quad (5.33)$$

where (b) follows from the existence of an associated eigenvector x such that $\rho(\mathbf{W})x \leq (d_i - 1) \cdot \tanh(J)$ with equality if and only if all \mathbf{C}_m have only non-zero values on the main diagonal, i.e., all variables have equal degree. As G_{N_2} has a larger portion of nodes with $d_i = \max d_i$ than G_{N_1} , the average degree is also higher; i.e., $\hat{d}_{G_{N_1}} < \hat{d}_{G_{N_2}}$. Then (a) follows from (5.32). \square

By combination of Theorem 3 and Theorem 4 we get:

Corollary 4.1. *The existence of a non-vanishing external field stabilizes BP on finite-size graphs.*

Next, we extend the above observations to varying couplings and fields. We assume that all couplings J_{ij} have the same sign, and all fields θ_i have the same sign, not necessarily the same as the couplings. Then the scaling coefficients c_k in (5.30) depend not only on d_i , but on J_{ij} and θ_i as well. Still, $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ can only contain either positive entries or negative entries; it follows that:

Corollary 4.2. *For infinite-size grid-graphs with either purely attractive interactions $J_{ij} > 0$ or purely repulsive interactions $J_{ij} < 0$, the existence of some non-vanishing fields $\theta_i \neq 0$ stabilizes BP.*

Theorem 5. *The eigenvalue-spectrum of the Jacobian $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ is symmetric if and only if the underlying graph is bipartite.*

Proof. The adjacency matrix of any bipartite graph can be rearranged and written in block form

$$\mathbf{A} = \begin{bmatrix} \mathbf{0} & \mathbf{M} \\ \mathbf{M}^T & \mathbf{0} \end{bmatrix}, \quad (5.34)$$

so that the eigenvalue-spectrum is symmetric [BH11]. By the same arguments as in Lemma 2 it follows that $\mathcal{F}'(\boldsymbol{\nu}^\circ)$ has the same structure as \mathbf{A} and thus a symmetric spectrum as well. \square

Corollary 5.1. *Assume we only have repulsive interactions. Then a graph is bipartite if and only if no frustrations occur.*

Proof. A cycle is frustrated if and only if the product of all J_{ij} along the corresponding edges is negative [MM09, p.45]. Frustrations can therefore only occur in graphs with cycles of odd length, which implies that the graph cannot be bipartite [KV05, Prop. 2.27]. \square

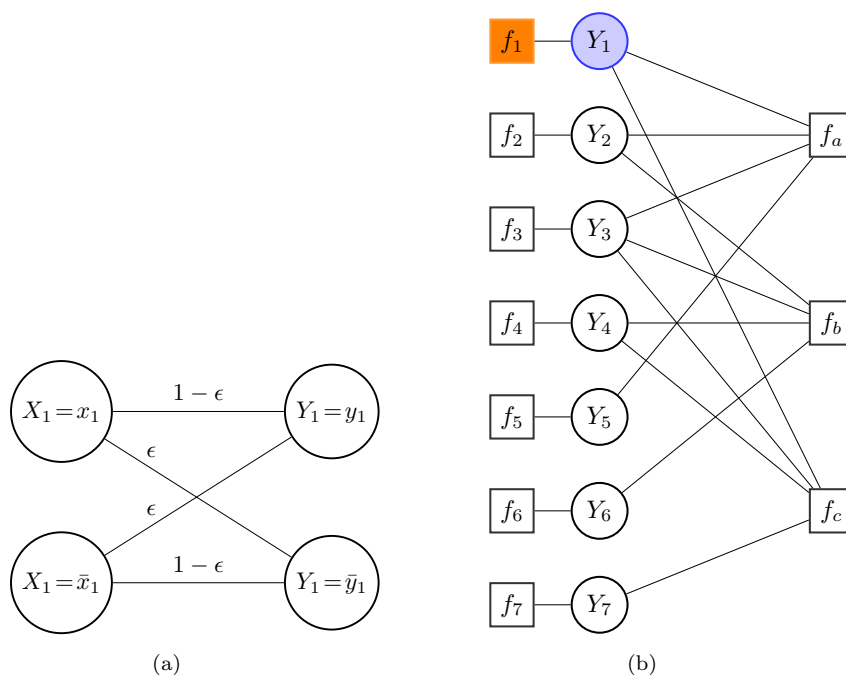


Figure 5.11: (a) Binary symmetric channel with error probability ϵ ; (b) factor graph for the (7,4) Hamming code that corresponds to (5.38) where Y_1 is flipped.

5.9 Application: Error-Correcting Codes

One of the most prominent applications where BP is successfully applied to loopy graphs is iterative decoding. We keep this section as self-contained as possible. For a thorough introduction we refer the interested reader to the textbooks [Mac03, Wym07]; the connection between BP and decoding is further explained in great detail in [KF98, AM00, KFL01].

We consider a binary symmetric channel (BSC) with a binary input $X_i \in \mathcal{X} = \{0, 1\} = \{x_i, \bar{x}_i\}$ and a binary output $Y_i \in \mathcal{Y} = \{0, 1\} = \{y_i, \bar{y}_i\}$. The channel is specified by the error-probability ϵ , where transmitted bits are flipped with probability ϵ . That is $P_{X_i|Y_i}(x_i|y_i) = P_{X_i|Y_i}(\bar{x}_i|\bar{y}_i) = 1 - \epsilon$ and $P_{X_i|Y_i}(x_i|\bar{y}_i) = P_{X_i|Y_i}(\bar{x}_i|y_i) = \epsilon$ (cf. Figure 5.11 (a)). Additional, redundant bits help to detect and correct transmission errors. The aim of error-correcting codes is to reach the desired error-correcting performance while introducing as little redundancy as necessary, i.e., to operate as close as possible to the theoretical limit. Suppose we transmit a codeword with block length $N = 7$ consisting of 4 source bits X_1, \dots, X_4 and three parity-check bits X_5, X_6, X_7 that satisfy

$$\begin{aligned} X_1 \oplus X_2 \oplus X_3 \oplus X_5 &= 0, \\ X_2 \oplus X_3 \oplus X_4 \oplus X_6 &= 0, \\ X_1 \oplus X_3 \oplus X_4 \oplus X_7 &= 0, \end{aligned}$$

where \oplus is an XOR, i.e., the sum in modulo-2 arithmetic. This linear irregular code is the (7,4) Hamming code [Mac03, Chapter 1].

In this example we assume that the sent message is $\mathbf{x} = (0, 0, 0, 0, 0, 0, 0)^{35}$ and that exactly one bit suffers from a bit flip. For irregular codes the degree of the variables d_i varies; therefore, we consider two scenarios: either $\mathbf{y} = (1, 0, 0, 0, 0, 0, 0)$ or $\mathbf{y} = (0, 0, 0, 0, 0, 1, 0)^{36}$.

It is often convenient to express a code in factorized form and represent it explicitly with a

³⁵ Note that the properties of the BSC are independent of the transmitted codeword \mathbf{x} .

³⁶ Normally the performance of a code is studied over an ensemble of sent codewords where each bit flips with probability ϵ .




Figure 5.12: Results for the $(7,4)$ Hamming code. We compare the exact solution $P_{X_i|Y}(x_i|y)$ (red) to the approximate solution $\tilde{P}_{X_i|Y}(x_i|y)$ of NPHC (blue) as ϵ increases for: (a) Y_1 is flipped and (b) Y_6 is flipped.

factor graph. A factor graph consists of variable nodes Y_i and factor nodes f_A where each factor f_A acts as a function on all variables connected $\mathbf{Y}_i = \{Y_i \in \partial(A)\}$.³⁷ On a factor graph, BP operates similar as introduced in Section 3.2. Now, two different types of messages are sent along every edge: factor-to-variable messages r_{Ai} and variable-to-factor messages q_{iA} . All messages are iteratively updated according to

$$r_{Ai}^{n+1}(y_i) = \sum_{\mathbf{y}_k: Y_k \in \{\partial(A) \setminus Y_i\}} f_A(\mathbf{Y}_k = \mathbf{y}_k, Y_i = y_i) \prod_{Y_k \in \{\partial(A) \setminus Y_i\}} q_{kA}^n(y_k), \quad (5.35)$$

$$q_{iA}^{n+1}(y_i) = \alpha_{iA}^n \prod_{f_B \in \{\partial(i) \setminus f_A\}} r_{Bi}^n(y_i), \quad (5.36)$$

where α_{iA} is chosen such that $q_{iA}^{n+1}(y_i) + q_{iA}^{n+1}(\bar{y}_i) = 1$. After all messages converged to a fixed point, the marginals of the variable nodes are approximated by the product of all incoming messages

$$\tilde{P}_{Y_i}(y_i) = \frac{1}{Z} \prod_{f_B \in \partial(i)} r_{Bi}^n(y_i). \quad (5.37)$$

Details of the factor graph representation can be found in [MM09, Chapter 9] and [Loe04].

Let us consider two types of factors: $f_i(Y_i) = P_{X_i|Y_i}(x_i|y_i)$ to model the BSC, and $f_A(\mathbf{Y}_i)$ to verify if all parity-checks are satisfied. Then $f_A(\mathbf{Y}_i) = 1$ if the sum of all arguments $\sum_{\mathbf{Y}_i} y_i$ is even and $f_A(\mathbf{Y}_i) = 0$ if the sum is odd. The conditional probability for $\mathbf{X} = \mathbf{x}$ to be the codeword, given the received codeword $\mathbf{Y} = \mathbf{y}$ then is

$$P_{\mathbf{X}|\mathbf{Y}}(\mathbf{x}|\mathbf{y}) = \frac{1}{Z} \prod_{i=1}^7 f_i(Y_i) \cdot f_a(Y_1, Y_2, Y_3, Y_5) \cdot f_b(Y_2, Y_3, Y_4, Y_6) \cdot f_c(Y_1, Y_3, Y_4, Y_7). \quad (5.38)$$

The corresponding factor graph representation is shown in Fig 5.11 (b).

Now we create a system of equations similar as in Section 5.3 and obtain all fixed points with NPHC. A unique fixed point exists for all settings – and this fixed point is stable, which justifies the application of BP on error-correcting codes. We further estimate the accuracy of

³⁷ Similar to Section 2.2, we use $\partial(\cdot)$ to specify the neighbors of nodes and variables.

the approximation; this relates to the following question. If we communicate over a BSC, how vulnerable is BP decoding to an increased error probability? To answer this question we obtain the fixed points for $\epsilon \in [0, 0.5]$ and compare the exact solution obtained by the junction tree algorithm $P_{X_i|\mathbf{Y}}(x_i|\mathbf{y})$ (red), to the approximate solution $\tilde{P}_{X_i|\mathbf{Y}}(x_i|\mathbf{y})$ obtained by NPHC (blue) in Figure 5.12.³⁸

An error can be corrected by BP decoding if $\tilde{P}_{Y_i|\mathbf{X}}(y_i|\mathbf{x}) > 0.5$. With exact decoding a single bit-flip can be corrected for $\epsilon < 0.21$. The fixed points obtained by NPHC, however, reveal that BP does not utilize the full potential of the code (Figure 5.12). If Y_1 was corrupted the error can be corrected for $\epsilon < 0.13$ (Figure 5.12 (a)). BP fails to correct the error if Y_6 was flipped for all values of ϵ (Figure 5.12 (b)). The reason therefore is a systematic error: the check-bit only has a single connection to the parity-check function f_c . According to (5.36) $q_{6,c}^{n+1}(y_i) = \alpha_{7,7} \cdot r_{7,7}(y_i)$; therefore Y_6 does not incorporate any information from the remaining graph. To conclude, the higher the connectivity of a node, the more information of other bits is taken into account and the better the error-correction capability of BP and NPHC.

5.10 Conclusion

This chapter was an attempt to get a deeper understanding of BP's behavior, with potential implications for deriving convergence guarantees and finding stronger conditions for uniqueness of BP fixed points.

The specific focus was to apply the tools from Chapter 4, namely the NPHC method and the concept of linearization to answer some of the open questions associated with BP (cf. Section 3.6). In particular we utilized the NPHC method to characterize the solution space and to obtain all BP fixed point solutions before performing a local stability analysis to understand the convergence properties.

Although this approach provided many interesting insights and answered some long-standing questions, some additional questions emerged in the process.

The focus on relatively small, well-structured models provides an obvious starting point and was essential in developing the results presented. Although many insights are expected to carry over to more general models, the generalization to larger models with fewer constraints on their potentials is primarily left open. Moreover, the general models considered here are seemingly too small as to exhibit an actual disordered behavior (characterized by the existence of multiple fixed points). The systems of polynomial equations considered so far, however, are already larger than most solvable problems; the generalization to even larger systems therefore remains problematic.

One key feature of the presented framework was the estimation of a favorable upper bound on the number of solutions. In particular, the BKK bound utilized the sparsity of the polynomial system induced by the graph structure and was tight in all our experiments. Computing the BKK bound and creating the start system, however, took up the major part of the overall runtime in our experiments. One will agree that the structure of the graph – together with the update equations – literally specifies the structure of the polynomial system. Yet, we only relied on generic, albeit well-established, methods to compute the BKK bound and to create the start-system. This suggests that one should take the graph structure into account, potentially yielding a much more efficient algorithm for computing the BKK bound.

Regarding the approximation quality of BP, we exemplified an accuracy-gap between fixed points obtained by BP and the best possible fixed points (obtained by NPHC). In practice this justifies the exploration of multiple fixed points: one can then either consider a combination

³⁸ Note that for $\epsilon = 0.5$ the transmission is random.

of weighted marginals, providing strikingly accurate results in the models considered; or one can select the fixed point that maximizes \mathcal{Z}_B , often leading to the best approximation of the marginals. We will further investigate this topic in Chapter 7 where we discuss the advantages and pitfalls of both approaches. Moreover, we analyzed the fixed point maximizing the partition function in detail: we revealed how it continuously deforms under varying parameters and how the stability depends on the graph structure. The observation of a continuous deformation will serve as the main inspiration for Chapter 6 where we build upon this observation to propose an enhanced version of BP.

Regarding the convergence properties of BP, we answered several questions on the basis of our empirical observations and our theoretical analysis. Graphs with vanishing local potentials are indeed a worst-case scenario; strong local potentials reduce the magnitude of all eigenvalues, thus helping to achieve convergence, and additionally increase the accuracy of the most accurate fixed point.

The graph structure affects the convergence properties as the spectral radius increases with the model size, which consequently degrades the performance of BP. Moreover, bipartite graphs exhibit a symmetric spectrum so that damping cannot help to achieve convergence; note that the ineffectiveness of damping for bipartite models was recently also observed for models with more general potentials [PI17]. Based on our observations it seems reasonable to conjecture that damping can only be used to stabilize a fixed point if it is unique. Whether this generalizes to spin glasses is by no means obvious. It would, for the purpose of understanding the convergence properties of spin glass models, be of interest to specify a distribution of the eigenvalues based on the distribution of the potentials.³⁹

³⁹ One could hope to obtain similar results as in the study of spectral densities on random matrices, where it is for example well-established that the eigenvalues of sparse symmetric matrices are distributed according to the Wigner semicircle law [Wig58, Erd11].

6

Enhancing Belief Propagation: Self-Guided Belief Propagation

“The most significant dimension of freedom is the freedom from one’s own ego - in other words, from the feeling that I am the center of everything.”
– Voytek Kurtyka

This chapter builds upon the insights from the preceding chapter and presents one possible way to make the performance of BP more robust. All models analyzed so far show that the minimum of the Bethe free energy is continuously deformed as the coupling strength increases. As a consequence of this observation, we progressively incorporate the pairwise potentials and keep track of the evolving fixed point. Although the idea of this algorithm was written down by the present author, such an approach may seem rather obvious; in fact Alexander Ihler revealed that he also had a similar idea in his mind already for quite some time [Ihl].

We begin this chapter with the introduction of our proposed algorithm *self-guided belief propagation* (SBP) in Section 6.2. In Section 6.3 we apply SBP to a wide range of models and discuss our empirical observations before we provide a formal analysis for the special case of attractive models with unidirectional local potentials in Section 6.4.

All theoretical results have been conducted by the present author. The implementation of the algorithm and the empirical evaluation in Section 6.3 were performed by the author’s student Florian Kulmer. Both the theoretical and empirical results have been prepared for publication in [KKP18].

6.1 Motivation

The analysis of BP’s solution space in Chapter 5 revealed that accurate fixed points may exist even though BP fails to obtain them. This is in accordance with the observation in [WJ14a] that the Bethe approximation is often accurate despite the failure of BP to converge. The main aim of this chapter lies in enhancing BP so that it converges to accurate fixed points while preserving its simple nature (of just considering local interactions).

Besides empirical analyses and some restricted theoretical studies, it remains an open problem to obtain a rigorous understanding of the limitations of BP for general graphs. Implementation details, as for example the initialization, play an important role in the case of multiple fixed points. In this case, it may depend on the initialization whether BP provides accurate marginals or not. The dependence of BP on such implementation details obviously poses a serious issue for providing performance guarantees.

One way to get convergence guarantees is to consider the equivalent optimization problem and to minimize the Bethe free energy \mathcal{F}_B . This, however, comes at the cost of an increased runtime complexity; polynomial-time algorithms only exist for restricted classes of problems and even approximating the global minimum might be problematic for graphical models with arbitrary potentials [CCG⁺11, Shi12, WJ14a]. Hence, the pursuit for methods that approximate the marginals with both runtime- and convergence-guarantees is still ongoing. These limitations

motivate the search for modifications of BP that overcome these issues in order to increase the accuracy and enhance the convergence properties. In this chapter, we introduce self-guided belief propagation (SBP) that aims to fill this gap.

The evolution of the fixed points revealed a close relationship between the coupling strength and the performance of BP. Especially strong couplings reduce the accuracy and deteriorate the convergence properties. This and the observation that tuning the coupling strength continuously deforms the fixed point minimizing the Bethe free energy inspired us to construct a homotopy.

More precisely, we first consider only local potentials (where BP is exact and has a unique fixed point) and subsequently modify the model by increasing the pairwise potentials to the desired values. SBP thus solves a deterministic sequence of models that iteratively refines the Bethe approximation towards an accurate solution that is uniquely defined by the initial model.

We evaluate SBP for grid-graphs, complete graphs, and random graphs with Ising potentials and, compared to BP, we observe superior performance in terms of accuracy; in fact SBP achieves more accurate results than Gibbs sampling in a fraction of runtime. We theoretically demonstrate optimality of the selected fixed point for *attractive* models with unidirectional local potentials. Additionally SBP enhances the convergence properties and excels for *general* models where SBP provides accurate results despite the non-convergence of BP. We further expect that the ease of use lowers the hurdle for practical applications.

6.2 Self-Guided Belief Propagation (SBP)

In this section we present an intuitive justification of the proposed method and subsequently introduce SBP in detail. We further present practical considerations and pseudocode of SBP. A formal treatment of SBP is deferred to Section 6.4.

The current understanding of BP is that strong (pairwise) potentials negatively influence BP and that in incorporating the potentials slowly [BKMZ07] may reduce the overall number of iterations. Inspired by our observations in Chapter 5 that strong local potentials increase accuracy and lead to better convergence properties, we aim to reduce the influence of the pairwise potentials that negatively influence BP.

SBP starts from a simple model with independent random variables and slowly incorporates the potential's strength, i.e., it solves the simple problem first and – by repetitive application of BP, keeps track of the fixed point as the interaction strength is increased by a scaling term. Again we resort to the homotopy continuation method for this purpose.

More formally, SBP considers an increasing length- K sequence $\{\zeta_k\}$ where $k = 1, \dots, K$ such that $\zeta_k < \zeta_{k+1}$ and $\zeta_k \in [0, 1]$ with $\zeta_1 = 0$ and $\zeta_K = 1$. This further indexes a sequence of probabilistic graphical models $\{\mathcal{U}_k\}$ that converges to the model of interest $\mathcal{U}_K = \mathcal{U}$. Every probabilistic graphical model has a set of potentials $\Psi_k = \{\Phi_k(x_i, x_j), \Phi_k(x_i)\}$ associated, where $\Phi_k(x_i) = \Phi_{X_i}(x_i)$ and the pairwise potentials at index k are exponentially scaled by

$$\begin{aligned} \Phi_k(x_i, x_j) &= \exp(J_{ij}\zeta_k x_i x_j) \\ &= \Phi_{X_i, X_j}(x_i, x_j)^{\zeta_k}. \end{aligned} \quad (6.1)$$

We further denote the fixed points of BP for \mathcal{U}_k by μ_k° . The initialization determines the performance of BP if multiple fixed points exist; SBP always provides a favorable initialization for the model \mathcal{U}_k by the preceding fixed point μ_{k-1}° and performs the composite function

$$\mu_K^\circ = \mathcal{BP}_K^\circ (\mathcal{BP}_{K-1}^\circ (\dots \mathcal{BP}_1^\circ (\mu_1^1))) . \quad (6.2)$$

This may lead to problems if the fixed point becomes unstable for some value $k < K$ in which case we cannot rely on BP to keep track of the fixed point anymore. Instead, SBP provides the last stable fixed point in that case, i.e., μ_{k-1}° as the final estimate.




Figure 6.1: Illustrative example: SBP proceeds along the smooth solution path and obtains accurate marginals despite instability of the terminal fixed point. Note that the fixed point becomes unstable for $\zeta > 0.7$; SBP stops and provides the last stable solution (that is already close to the exact one) as an approximation.

In other words, SBP relaxes the problem of minimizing $\mathcal{F}_{\mathcal{B}}$ by making all variables independent (and the Bethe approximation exact). Then, the problem is deformed into the original one by increasing ζ from zero to one. Thereby, a stationary point $\mathcal{F}_{\mathcal{B}}^{\circ}$ emerges as a well-behaved path (cf. Proposition 1 in Section 6.4) and SBP keeps track of it with BP constantly correcting the stationary point.

We illustrate how SBP approximates the marginals for a problem where BP fails to converge in Figure 6.1. Initially, SBP obtains the pseudomarginals for $\zeta = 0$ and then estimates the marginals of the desired problem by successively increasing ζ and running BP to keep track of the emerging solution path. Note that the approximated marginals are already close to the exact ones in this example; experiments show that this is often the case (cf. Section 6.3).

6.2.1 Practical Considerations

The procedure of SBP is essentially a path-tracking problem. As already discussed for the NPHC method in Section 4.3.3, such a problem consists of three key steps: prediction, correction, and step-size adaption. While the implementation details of these steps may influence the performance, we have found in our experiments that the efficiency of SBP remains largely unaffected. We will now describe the specific details used for all three steps throughout our experiments.

In practice the runtime of SBP is influenced by the difference between two successive fixed points μ_k° and μ_{k-1}° – the difference is primarily determined by the number of steps K . Ideally, to reduce the difference, K should be as large as possible. This, however, increases the runtime as well (cf. Theorem 8); in practice we would choose K as small as possible but as large as necessary. Moreover, one can adaptively increase the step size if two successive fixed points are close, i.e., if $\mu_k^{\circ} \simeq \mu_{k-1}^{\circ}$ (cf. [SW05, pp.23], [AG03]). Our experiments show that it is sufficient to use rather coarse steps (we used $K \leq 10$ for all reported experiments).

Additionally, instead of initializing \mathcal{BP}_k with its preceding fixed point messages, i.e., $\mu_k^1 = \mu_{k-1}^{\circ}$ one can (e.g., by spline extrapolation) estimate $\mu_k^1 = f(\mu_{k-1}^{\circ}, \mu_{k-2}^{\circ}, \dots, \mu_{k-l}^{\circ})$ so that $\mu_k^1 \cong \mu_k^{\circ}$ to reduce the overall number of iterations. We empirically observed that the benefit diminishes for $l > 3$.

6.2.2 Pseudocode

Pseudocode of SBP is presented in Algorithm 1. The maximum number of iterations for BP is given by $N_{BP} = 10^3$. We randomly initialize μ_1^1 and either use fixed step size or adaptive step

size (*adaptive stepsize* = 1). The sequences of messages is contained in $\{\mu_k^\circ\} = \{\mu_1^\circ, \dots, \mu_k^\circ\}$.

Cubic spline extrapolation is applied in `ExtrapolateMsg` to estimate the initial messages of the subsequent model.

We further present the pseudocode for the adaptive step size controller in Algorithm 2.

Algorithm 1: Self-Guided Belief Propagation (SBP)

input : Graph $\mathcal{G} = (\mathbf{X}, \mathbf{E})$, Potentials Ψ
output: Fixed point messages μ°

- 1 initialization $\mu_1^1 \leftarrow \mu^1$
- 2 $k \leftarrow 1$
- 3 $step_{init} \leftarrow \frac{1}{10}$
- 4 $\zeta_1 \leftarrow 0$
- 5 **while** $\zeta \leq 1$ **do**
- 6 $\Psi(\zeta_k) \leftarrow \text{ScalePotentials}(\Psi, \zeta_k)$
- 7 $(\mu, n) \leftarrow \text{BP}(\mu_k^1, \Psi(\zeta_k), N_{BP})$
- 8 **if** $n < N_{BP}$ **then**
- 9 $\mu_k^\circ \leftarrow \mu$
- 10 **else**
- 11 **break**
- 12 **if** *adaptive stepsize* **then**
- 13 $\zeta_{k+1} \leftarrow \zeta_k + \text{AdaptiveStepSize}(\{\mu_k^\circ\}, step_{init}, k)$
- 14 **else**
- 15 $\zeta_{k+1} \leftarrow \zeta_k + step_{init}$
- 16 $\mu_{k+1}^1 \leftarrow \text{ExtrapolateMsg}(\{\mu_k^\circ\}, \{\zeta_k\})$
- 17 $k \leftarrow k + 1$
- 18 $\mu^\circ \leftarrow \mu_{k-1}^\circ$

Algorithm 2: Adaptive Step Size Controller

input : Sequence of messages $\{\mu_k^\circ\}$, $step_{init}$, k
output: $step$

- 1 $step \leftarrow step_{init}$
- 2 $threshold \leftarrow 1 \cdot 10^{-3}$
- 3 $l \leftarrow 1$
- 4 **while** $(\text{MSE}(\mu_k^\circ) - \text{MSE}(\mu_{k-l}^\circ)) < threshold$ **do**
- 5 $l \leftarrow l + 1$
- 6 $step \leftarrow step + step_{init} \cdot l$

6.3 Experiments

We apply SBP to attractive (Section 6.3.2) and general (Section 6.3.3) models on $n \times n$ grid graphs of different size, and to complete and random graphs (average degree of $\hat{d}_G = 3$) with $N = 10$ random variables. Experiments were performed for these graphs in order to render the computation of the exact marginals feasible and to make the results comparable to previous work [WTJS14, SJ08, MJGF09, SRF16].

6.3.1 Experimental Settings

SBP is evaluated and compared to BP, BP_D (BP with damping), and Gibbs sampling. The exact marginals are obtained by the junction tree algorithm and the accuracy of the marginals is evaluated by the mean squared error $E_P(m)$ between the approximate marginals at the m^{th} fixed point $\tilde{P}_{X_i}^m(x_i)$ and the exact marginals $P_{X_i}(x_i)$. Additionally, we approximate the global minimum \mathcal{F}_B^* by [WJ14a] and evaluate the mean squared error (MSE_B) between the approximate marginals and the marginals obtained at the global minimum of the Bethe free energy $\tilde{P}_{X_i}^*(x_i)$. We further compare the runtime of all methods by counting the overall number of BP iterations and the number of iterations for Gibbs sampling.⁴⁰

We consider $L = 100$ models with random potentials for every experiment. The initial messages are randomly initialized 100 times for each of these L models, before applying BP with and without damping. We consider BP (and BP_D) as converged if at least a single message initialization (out of 100) exists for which BP converges. We report the convergence ratio, i.e., the number of experiments (or probabilistic graphical models) for which BP converged at least once divided by the overall number of models L .

The reported mean-squared error $E_P(m)$ and the number of iterations are averaged over all convergent runs of BP and BP_D (i.e., BP^o and BP_D^o) while all runs that did not converge are discarded. SBP, on the other hand, allows obtaining an approximation of the terminal fixed point in case that this fixed point is unstable, which prevents BP and BP_D from converging. Therefore, we average the error and the number of iterations over all L models for SBP (SBP_{all}), Gibbs sampling (Gibbs_{all}), and for minimization of the Bethe approximation (\mathcal{F}_B^*).

For BP and SBP we set the maximum number of iterations to $N_{BP} = 10^3$ and use random message scheduling. For BP_D we choose a large damping factor $\epsilon = 0.9$ to account for strong couplings. Such a large damping factor helps to prioritize convergence over runtime – this admits comparison of marginal accuracy for a wide range of models. The maximum number of iterations, however, has to be increased to $N_{BP} = 10^4$ in order to account for the slower convergence. Carefully selecting a damping factor that depends on a given model may reduce the number of iterations until convergence but can not increase the accuracy; moreover, if chosen too small BP_D may fail to converge. The accuracy of SBP is only marginally affected by its parameters and we use the following parameters for all experiments: $K \leq 10$, adaptive step size, and cubic spline extrapolation. Gibbs sampling is run for 10^5 iterations.

⁴⁰ Computing the acceptance-probability requires similar runtime as one BP message update.




Figure 6.2: MSE, i.e., $E_P(m)$, and number of iterations for: SBP_{all} (blue), BP^o (magenta), and BP_D^o (green); the local potentials are specified by $\theta_i \sim \mathcal{U}(-0.5, 0.5)$ and the pairwise potentials are specified by (a) $J_{ij} \sim \mathcal{U}(0, \beta)$ (attractive model); (b) $J_{ij} \sim \mathcal{U}(-\beta, \beta)$ (general model).

6.3.2 Attractive Models

We consider grid graphs with $N = 100$ random variables (10×10), random graphs with $N = 10$ random variables, and complete graphs with $N = 10$ random variables. For each graph, we generate $L = 100$ models for every value of $\beta \in \{0, 0.5, \dots, 5\}$ and sample the potentials according to $\theta_i \sim \mathcal{U}(-0.5, 0.5)$ and $J_{ij} \sim \mathcal{U}(0, \beta)$; i.e., we consider 1100 different parametrizations. Note that BP is randomly initialized 100 times for every considered parametrization. We compute $E_P(m)$ for every value of β and visualize the mean and the standard deviation of the MSE⁴¹ as well as the number of iterations in Figure 6.2 (a).

Note that BP (magenta) converges rapidly for all graphs considered; hence, there is no additional benefit for BP_D (green) that only increases the number of iterations. SBP (blue) only slightly increases the number of iterations as compared to BP and converges in fewer iterations than BP_D. Note that SBP is guaranteed to capture the global optimum if all local potentials are unidirectional (cf. Theorem 9-10). But even if we do allow for random local potentials, we em-

⁴¹ Note that the mean-squared error is not Gaussian distributed but we report the standard deviation for simplicity.

Table 6.1: RESULTS FOR GENERAL MODELS WITH $J_{ij} \in \{-1, 1\}$ ON GRID GRAPHS ($N = 25$ AND $N = 100$), COMPLETE GRAPHS ($N = 10$), AND RANDOM GRAPHS ($N = 10$). WE REPORT THE MSE TO THE EXACT MARGINALS, I.E., $E_P(m)$, AND THE MSE TO THE BETHE APPROXIMATION, I.E., $\text{MSE}_{\mathcal{B}}$, CONVERGENCE RATIO, AND THE OVERALL NUMBER OF BP ITERATIONS. ONLY CONVERGED RUNS ARE CONSIDERED FOR BP° AND BP_D° BUT ALL RUNS ARE CONSIDERED FOR SBP_{all} , GIBBS_{all} , AND $\mathcal{F}_{\mathcal{B}_{all}}^*$.

	θ	Grid Graph (5×5)			Grid Graph (10×10)			Complete Graph			Random Graph		
		0	0.1	0.4	0	0.1	0.4	0	0.1	0.4	0	0.1	0.4
$E_P(m)$	BP°	0.338	0.251	0.102	-	-	0.184	0.463	0.466	0.356	0.252	0.202	0.101
	BP_D°	0.226	0.198	0.066	0.186	0.240	0.154	0.463	0.473	0.422	0.128	0.116	0.083
	SBP_{all}	0.000	0.029	0.047	0.000	0.026	0.077	0.000	0.055	0.074	0.000	0.048	0.049
	$\mathcal{F}_{\mathcal{B}_{all}}^*$	0.036	0.042	0.069	-	-	-	-	-	-	-	-	-
	Gibbs_{all}	0.001	0.016	0.064	0.001	0.037	0.120	0.096	0.096	0.077	0.001	0.011	0.048
Convergence ratio	BP°	0.05	0.11	0.26	0.00	0.00	0.02	0.41	0.42	0.50	0.30	0.33	0.49
	BP_D°	0.11	0.16	0.69	0.01	0.02	0.12	0.41	0.41	0.50	0.62	0.64	0.80
Number of iterations	BP°	40	52	84	-	-	102	17	17	18	42	53	50
	BP_D°	1370	1449	1735	2711	2313	2599	211	207	234	1077	1057	873
	SBP_{all}	5	182	146	5	149	209	5	51	110	5	149	131
	Gibbs_{all}	10^5	10^5	10^5	10^5	10^5	10^5	10^5	10^5	10^5	10^5	10^5	10^5
MSE $_{\mathcal{B}}$	SBP_{all}	0.036	0.037	0.022	-	-	-	-	-	-	-	-	-
Global Min	SBP	100	10	23	-	-	-	-	-	-	-	-	-

empirically observe that SBP consistently outperforms BP with respect to accuracy. This becomes especially evident for models with strong couplings: these models exhibit multiple stable fixed points [KMCP18] such that, depending on the initialization, BP often converges to inaccurate fixed points.

6.3.3 General Models

General models admit frustrated cycles and traditionally pose problems for BP and other methods that aim to minimize the Bethe approximation.

First, in order to evaluate the performance of SBP we consider $\theta_i = \theta \in \{0, 0.1, 0.4\}$ and draw the couplings with equal probability from $J_{ij} \in \{-1, 1\}$; the results are summarized in Table 6.1. Although BP and BP_D° fail to converge for most models we observe that SBP stops after only a few iterations and significantly outperforms BP in terms of accuracy. In fact, SBP achieves accuracy competitive with Gibbs sampling but requires three orders of magnitude fewer iterations.

Second, we further apply SBP to general graphs and evaluate whether SBP provides a good approximation of the pseudomarginals that correspond to the global minimum of the Bethe free energy according to $\tilde{P}_B^* = \arg \min_{\mathbb{L}} \mathcal{F}_{\mathcal{B}}(\tilde{P}_B)$ (cf. (3.29)). Therefore we consider grid graphs (of size 5×5), which still allows us to approximate $\mathcal{F}_{\mathcal{B}}^*$ – and the related pseudomarginals – reasonable well by [WJ14a]. The results are summarized in Table 6.1 and show that SBP approximates \tilde{P}_B^* within the accuracy of our reference method ($\text{MSE}_{\mathcal{B}}$). We further report the number of times where SBP obtains the terminal fixed point, i.e., for \mathcal{U}_K , in Table 6.1 (Global Min). It becomes obvious that SBP approximates the terminal fixed point reasonably well, despite frequently stopping for $\zeta_k < 1$. Moreover, closer inspection of the accuracy reveals that SBP does not only approximate the “correct” pseudomarginals well ($\text{MSE}_{\mathcal{B}}$), but concurrently provides an accurate approximation of the exact marginals (MSE).

Third, we investigate how the approximation quality depends on the scaling parameter ζ_k . Therefore, we depict the evolution of the MSE (to the exact solution) and $\text{MSE}_{\mathcal{B}}$ (to the approximate solution) in Figure 6.3. We observe that $\text{MSE}_{\mathcal{B}}$ (blue) decreases monotonically with every iteration, which empirically verifies that SBP proceeds along a well-behaved solution path




Figure 6.3: MSE, i.e., $E_P(m)$, (orange) and MSE_B (blue) over the cumulative number of iterations. Results are averaged over 100 grid graphs (5×5) with $\theta_i = 0.4$ and $J_{ij} \in \{-1, 1\}$.

(cf. Proposition 1). Note that MSE_B decreases rapidly in the first iterations and SBP spends a major part of the overall runtime for slight improvements. The MSE to the exact solution, on the other hand, decreases first until it increases again as SBP incorporates stronger couplings. Stronger couplings tend to degrade the quality of the Bethe approximation in loopy graphs and lead to marginals that are increasingly biased towards one state [Wei00, KMCP18]. This explains why the MSE to the exact solution increases as SBP converges towards the terminal fixed point. One could exploit this behavior and restrict the runtime by stopping SBP after consumption of a fixed iteration budget; this may even increase the accuracy with respect to the exact solution.

Finally, we aim to investigate the influence of the coupling strength: therefore we consider $\theta_i \sim \mathcal{U}(-0.5, 0.5)$ and $J_{ij} \sim \mathcal{U}(-\beta, \beta)$. For every $\beta \in [0, 5]$ we execute $L = 100$ experiments and present the averaged results in Figure 6.2 (b). Note that we restrict the results to $\beta \leq 2$ on the grid graph because BP did only converge sporadically for models with stronger couplings. SBP requires only slightly more iterations than BP and fewer than BP_D , even though we compare only to models where BP (or BP_D) converged. The benefits of SBP become increasingly evident as the coupling strength increases. Again SBP (blue) significantly outperforms BP° (magenta) and BP_D (green) on all graphs with respect to accuracy.

6.4 Theoretical Analysis of Self-Guided Belief Propagation

Here we present some more formal arguments and discuss the properties of SBP to understand under which conditions the algorithm (presented in Section 6.2) can be expected to perform well. We only present the most important Theorems and their implications below and defer the longer proofs to Appendix A.1.

6.4.1 Definitions

First, we fix our notation: we denote the pseudomarginals, corresponding to local minima of the Bethe free energy, of \mathcal{U}_k by $\tilde{P}_B^m(\zeta_k)$, and, with slight abuse of notation, we refer to the corresponding stationary point of the Bethe free energy by $\mathcal{F}_B^m(\zeta_k) = \mathcal{F}_B(\tilde{P}_B^m(\zeta_k))$. Note that the superscript m accounts for the fact that we are only looking at the local minima of \mathcal{F}_B (cf. Section 3.3.2).

It is beneficial to study the behavior of SBP as K tends towards infinity. Therefore we consider the unit interval $\zeta \in [0, 1]$ to be the compact support of the functions $\mathcal{F}_B(\zeta)$ and $\tilde{P}_B(\zeta)$. SBP is




Figure 6.4: Solution path (cf. Section 6.4.1) for a grid graph; start- and endpoint are depicted by blue points. Note how Propositions 1.1 - 1.2 are fulfilled.

inspired by the idea to proceed along a so-called *solution path* as ζ increases from zero to one in order to obtain the marginal distributions for the model of interest. Therefore, we shall consider a continuous homotopy function $\mathbf{H}(\boldsymbol{\mu}, \zeta) : \mathbb{R}^{|\boldsymbol{\mu}|+1} \rightarrow \mathbb{R}^{|\boldsymbol{\mu}|}$ that is defined by

$$\mathbf{H}(\boldsymbol{\mu}, \zeta) = \boldsymbol{\mu} - \mathcal{BP}(\boldsymbol{\mu}) \quad \text{where} \quad \Psi = \Psi(\zeta). \quad (6.3)$$

Then, a *solution path*

$$c(\zeta) : \mathbf{H}(\boldsymbol{\mu}, \zeta) = 0 \quad (6.4)$$

exists that (i) has a start point $c(\zeta = 0) = \boldsymbol{\mu} : \mathbf{H}(\boldsymbol{\mu}, \zeta = 0) = 0$, (ii) an endpoint $c(\zeta = 1) = \boldsymbol{\mu} : \mathbf{H}(\boldsymbol{\mu}, \zeta = 1) = 0$, and (iii) is continuous over $\zeta \in [0, 1]$, i.e., it connects the start- with the endpoint. SBP then proceeds along some solution path from a given start- to its endpoint. Note that the solution path $c(\zeta)$ is defined along the solutions to the fixed point equations and thus it implicitly defines the pseudomarginals $\tilde{P}_B^m(\zeta)$ by (3.2) and (3.3). In particular, we refer to the start- and endpoint by $\tilde{P}_B^m(\zeta = 0)$ and $\tilde{P}_B^m(\zeta = 1)$ respectively.

The following example in Figure 6.4 illustrates the fixed point evolution for a grid graph with attractive couplings. This example exhibits a unique solution path according to our definition; note, however, that a second curve exists, which lacks a start point and is therefore of no relevance for any method that proceeds along a solution path defined by the homotopy in (6.3).

6.4.2 Properties of Self-Guided Belief Propagation

The following proposition summarizes the main properties of the solution path that is specified and followed by SBP.

Proposition 1 (Properties for attractive and general models).

- (1) BP has a unique fixed point $\boldsymbol{\mu}_1^\circ$ for $\zeta_1 = 0$, so that SBP has a unique start point $\tilde{P}_B^m(\zeta = 0)$. (cf. Theorem 6)
- (2) A smooth (i.e., continuous) solution path originates from the start point $\tilde{P}_B^m(\zeta = 0)$. (cf. Theorem 7)
- (3) SBP efficiently proceeds along this (unique) solution path. (cf. Theorem 8)

Theorem 6 (Proposition 1.1). *A unique solution exists for $\zeta = 0$, i.e., a single start point $\tilde{P}_B^m(\zeta = 0)$ exists, and BP is guaranteed to converge. Moreover, this start point is exact, i.e., $\tilde{P}_{X_i}(\zeta = 0) = P_{X_i}(\zeta = 0)$ for all $X_i \in \mathbf{X}$.*

Proof. First, let us obtain the singleton marginals by summing over the pairwise marginals such that $\tilde{P}_{X_i}(x_i) = \sum_{x_j \in \mathcal{X}} \tilde{P}_{X_i, X_j}(x_i, x_j)$. For $\zeta_1 = 0$, that is for $\Phi_1(x_i, x_j) = 1$ it follows that marginalizing over the pairwise marginals (cf. (3.3)) equates to

$$\tilde{P}_{X_i}(x_i) = \Phi(x_i) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^\circ(x_i) \cdot \sum_{x_j \in \mathcal{X}} \Phi(x_j) \prod_{X_l \in \{\partial(j) \setminus X_i\}} \mu_{lj}^\circ(x_j). \quad (6.5)$$

Note that according to (3.1) $\mu_{ji}^\circ(x_i)$ is equivalent to the second part of (6.5) so that

$$\tilde{P}_{X_i}(x_i) = \Phi(x_i) \prod_{X_k \in \partial(i)} \mu_{ki}^\circ(x_i), \quad (6.6)$$

which equals (3.2). It follows that $\tilde{P}_{X_i}(+1) = e^{\theta_i} / (e^{\theta_i} + e^{-\theta_i}) = P_{X_i}(+1)$. \square

Note that Theorem 6 concurs with the sandwich-bound [WJ13, Th.4] that reduces to $P_{X_i}(+1) = \theta_i / (e^{\theta_i} + e^{-\theta_i}) = \tilde{P}_{X_i}(+1)$ for $J_{ij} = 0$. Theorem 6 thus reduces the problem of initializing SBP to computing μ_1° , which can be done in linear time.

Theorem 7 (Proposition 1.2). *Let $\tilde{P}_B^m(\zeta)$ be the pseudomarginals that are defined along the unique solution path that originates from $\tilde{P}_B^m(\zeta = 0)$. Then, $\tilde{P}_B^m(\zeta)$ and the associated stationary points $\mathcal{F}_B^m(\zeta)$ are continuous on their compact support $\zeta \in [0, 1]$.*

Proof. First, we show that the Bethe free energy $\mathcal{F}_B(\zeta)$ itself is an analytic function. Consider (3.24) with the pairwise potentials defined by (6.1). Then, the derivative with respect to ζ is given by

$$\begin{aligned} \frac{\partial \mathcal{F}_B(\zeta)}{\partial \zeta} &= -\frac{\partial}{\partial \zeta} \sum_{(i,j) \in \mathbf{E}} \sum_{x_i, x_j} \tilde{P}_{X_i, X_j}(x_i, x_j) \ln \Phi(x_i, x_j) \\ &= -\frac{\partial}{\partial \zeta} \sum_{(i,j) \in \mathbf{E}} \sum_{x_i, x_j} \tilde{P}_{X_i, X_j}(x_i, x_j) \cdot \zeta \cdot J_{ij} \cdot x_i x_j \\ &= -\sum_{(i,j) \in \mathbf{E}} J_{ij} \cdot \chi_{ij}. \end{aligned} \quad (6.7)$$

As an immediate consequence, we observe that $\mathcal{F}_B(\zeta)$ is continuously differentiable⁴² as (6.7) is a finite sum over finite terms.⁴³

We specifically consider the minimum $\mathcal{F}_B^m(\zeta)$ that emerges from the global minimum $\mathcal{F}_B^m(\zeta = 0) = \mathcal{F}_B^*(\zeta = 0)$; this start point is unique by Theorem 6. It follows by (6.7) that $\mathcal{F}_B^m(\zeta)$ varies in a continuous fashion along the unique solution path for $\zeta \in [0, 1]$. Further note that stationary points are in a one-to-one correspondence with fixed points of BP which completes the proof.

Further note that the set of stationary points is finite (cf. Section 3.3.1)⁴⁴ and that pitchfork bifurcations may only occur if $\theta_i = 0$ [PAM11], in which case SBP obtains the exact solution (cf. Theorem 9). \square

Theorem 7 substantiates the claim that a smooth solution path emerges from the simple problem.

⁴² Strictly speaking $\mathcal{F}_B(\zeta)$ is an analytic function.

⁴³ This is in accordance with the fact that true phase transitions (singularities in the derivative of the free energy) can occur only in the thermodynamic limit, where (6.7) is an infinite sum that equates to infinity.

⁴⁴ This is also required for the one-step replica symmetry breaking assumption [MM09, Section19].

Theorem 8 (Proposition 1.3). *There exists some scaling factor $\zeta_k \leq 1$ for which SBP converges to $\tilde{P}_B^m(\zeta_k)$ in $\mathcal{O}(KN_{BP})$.*

Proof of Theorem 8. SBP increases ζ_k as long as BP converges in fewer than N_{BP} iterations, and stops otherwise. Consequently, BP corrects the accuracy of the fixed point for each value ζ_k within a bounded number of iterations. The runtime of SBP is further determined by the choice of K , i.e., the step-size (cf. Section 6.2.1). Assume that SBP converges for ζ_k , then it does so in $\mathcal{O}(K \cdot N_{BP})$. \square

SBP is consequently capable of efficiently tracking the fixed point that emerges as ζ increases and requires KN_{BP} iterations at most. SBP may, however, only converge to a surrogate model for $\zeta_k < 1$ and is not guaranteed to obtain the pseudomarginals of the desired problem. One can characterize this error by computing a bound on $|\mathcal{F}_B^m(\zeta_k) - \mathcal{F}_B^m(\zeta_K)|$ given the difference between Ψ_k and Ψ_K (cf. [IFW05, Theorem 16]).

Corollary 8.1. *SBP obtains the pseudomarginals of the desired problem if and only if the endpoint $\tilde{P}_B^m(\zeta = 1)$ is stable, i.e., if $\tilde{P}_B^m(\zeta = 1) \in \mathbf{S}$. This is an immediate consequence of the fact that the convergence properties can only degrade along a given solution path (cf. Section 5.8).*

Theorem 6-8 are of fundamental importance but do not relate to the accuracy of the obtained stationary point. Assessing the quality of the Bethe approximation and the accuracy of BP for general models is still an open research question that is beyond the scope of this thesis. However, we further present Proposition 2 to discuss the accuracy of the obtained solution for *attractive* models.

Proposition 2 (Properties for attractive models with unidirectional fields). *The solution path $c(\zeta)$ leads towards an accurate solution with $\tilde{P}_B^m(\zeta_k) = \tilde{P}_B^*(\zeta_k)$ and $\mathcal{F}_B(\tilde{P}_B^m(\zeta_k)) = \mathcal{F}_B(\tilde{P}_B^*(\zeta_k))$. (cf. Theorem 9 and Theorem 10)*

We start by generalizing Griffiths' inequality [Gri67] to the fixed points of BP (Lemma 4) and subsequently provide Theorem 9-10 that discuss the accuracy of the fixed point obtained by SBP.

Lemma 4. *Consider two attractive probabilistic graphical models \mathcal{U}_0 and \mathcal{U}_1 with equal \mathcal{G} and with the potentials specified by $\theta_i > 0$ and by J_{ij}^0 and J_{ij}^1 , where $J_{ij}^0 < J_{ij}^1$ for all $(i, j) \in \mathbf{E}$. Let us consider a fixed point of BP with positive means $m_i^0 \in (0, 1]$. Then, $m_i^0 < m_i^1$ and $\chi_{ij}^0 < \chi_{ij}^1$.*

The proof of Lemma 4 contains some tedious, although not too complicated, algebraic manipulations and is thus deferred to Appendix A.1.1

Theorem 9 (Prop. 2). *Consider an attractive model with $\theta_i > 0$.⁴⁵ Then, $m_i(\zeta)$ increases monotonically along the solution path $c(\zeta)$; in particular SBP minimizes the Bethe approximation error and is optimal with respect to marginal accuracy, i.e.,*

$$\begin{aligned} \tilde{P}_B^m(\zeta) &= \arg \min_{\tilde{P}_B^k \in \mathbf{M}} E_P(k) \\ &= \arg \min_{\tilde{P}_B^k \in \mathbf{M}} E_Z(k). \end{aligned} \tag{6.8}$$

We will only sketch the proof of Theorem 9 here and refer to Appendix A.1.2 for a more elaborate treatment. Essentially, we have to bring the preceding Theorems together to show that SBP obtains the most accurate marginals for attractive models with unidirectional fields.

Theorem 6 explains how the initial model can be solved exactly and that it has positive mean and correlation for $\theta_i > 0$. Consequently, Lemma 4 and Theorem 7 apply and guarantee a

⁴⁵ Note that equal results can be obtained for $\theta_i < 0$ because of symmetry properties.

continuous solution path with monotonically increasing means $m_i(\zeta)$. Optimality of the obtained fixed point then is an immediate consequence of the definition of $\mathcal{F}_{\mathcal{B}}$ (cf. (3.24)), the RSB assumption (which holds for these models, cf. Section 5.6.3), and the existence of at most two local minima, i.e., $\mathbf{M} \leq 2$, for the considered models (cf. Lemma 5 and [WTJS14]).

Note that an attractive model with unidirectional fields is a special case of a (larger) attractive model with vanishing fields (cf. [Fis67] and [SKZ17, Section 1.2]).

Theorem 10 (Proposition 2). *Consider an attractive model with $\theta_i = 0$. Then, SBP obtains the exact solution, i.e., $\tilde{P}_{\mathcal{B}}^m(\zeta = 1) = P_{\mathcal{B}}(\zeta = 1)$.*

Proof. Only considering attractive models makes it straightforward to calculate the exact solution if all $\theta_i = 0$. The exact solution always has zero mean for all random variables, i.e., $m_i(\zeta) = 0$ for all values of ζ . Moreover, the exact solution coincides with a stationary point $\mathcal{F}_{\mathcal{B}}^m(\zeta_k)$ [MK07]. Depending on the coupling strength, this stationary point is a stable minimum (for sufficiently small values of $J_{ij} < J_C$) and a local maximum (for $J_{ij} > J_C$) (cf. Section 5.7.1 or to [MM09, pp.385]).

SBP consistently obtains this fixed point irrespective of its stability. Therefore, note that by Theorem 6 all $m_i(\zeta = 0) = 0$ and all messages are equal, i.e., $\mu_1^q(\zeta = 0) = 1/2$ for all $(i, j) \in \mathbf{E}$. SBP remains exactly on this fixed point and obtains the exact marginals as these fixed point messages can be represented exactly, without any quantization errors in binary arithmetic. \square

To conclude, for attractive models with $\theta_i \leq 0$ or $\theta_i \geq 0$, SBP either obtains the fixed point that corresponds to the global minimum $\mathcal{F}_{\mathcal{B}}^*$ (Theorem 9), or it obtains the fixed point that corresponds to the exact solution, i.e., to $\min \mathcal{F}_{\mathcal{G}}$ (Theorem 10).

For *general models* m_i , need not increase monotonically along the solution path and it is not obvious whether SBP obtains the most accurate fixed point. However, the experimental results in Section 6.3.3 at least corroborate that SBP converges to accurate fixed points in many cases.

6.5 Conclusion

This chapter aimed at deriving a simple and robust method to efficiently perform approximate inference for models where BP fails to converge. The observations from our solution space analysis in Chapter 5 suggested that an accurate fixed point lies at the end of a solution path that emerges from the origin (where $J_{ij} = 0$). We managed to enforce convergence towards this fixed point by changing the updated function accordingly. To do so, we resorted to the homotopy method once again and utilized its proven capability of keeping track of solution paths, where we relied on BP in the correction step.

Our theoretical analysis of SBP for attractive models with unidirectional local fields revealed the existence of a unique, well-behaved solution path that leads to the best possible fixed point of BP and that can be tracked successfully by SBP. We further applied SBP to general models, where, despite the lack of theoretical guarantees, it exhibited a promising performance. Overall SBP consistently obtained marginals of better quality than BP with and without damping and approximated the marginals well on models for which BP did fail to converge at all.

One surprising effect of SBP was that it sometimes provided marginals that were even more accurate than the ones at the global minimum of $\mathcal{F}_{\mathcal{B}}$. This occurred frequently, specifically whenever SBP terminated early because of BP failing to converge in the correction step before the pairwise potentials were set to their desired value.

So far we can only provide some hand-waving arguments for this behavior: Therefore, note that common wisdom of BP tells us that BP usually fails to converge because of strong pairwise potentials; additionally, we have seen how the marginals become overconfident and tend to be

strongly biased towards one state as the coupling-strength increases. The latter property explains why the approximation quality of \mathcal{F}_B degrades as its minimum moves further away from the minimum of \mathcal{F}_G . This implies that following the solution path to its very end only makes the accuracy of the estimated marginals worse. It seems that the former property implicitly prevents this from happening and that the stability breaks down somewhat close to the exact solution.

First empirical evaluations suggest that this is indeed quite often the case. This raises an intriguing theoretical question: Is there really some fundamental connection between the exact marginals and the approximated marginals that are obtained at the onset of instability?– and if so, what is the nature of this connection?

In a more concrete manner, it also remains to generalize our theoretical analysis. So far we have only discussed the case of attractive models with unidirectional local fields, which – admittedly – is a rather restrictive setting. Although this restriction simplified our analysis significantly and provided some interesting insights into the nature of SBP, it still leaves some room for improvement.

Going back to our theoretical analysis once more, our statement of optimality relied on the fact that the global minimum of \mathcal{F}_B is optimal with respect to the accuracy of both the marginals and the free energy. Yet, it is not obvious if such a relationship generalizes to more general models as well. We will thus devote the next chapter to studying the underlying principles that relate the approximation quality of the marginals and of the free energy to each other.

7

Understanding Belief Propagation: Accurate Marginals or Partition Function

*"To the wise, life is a problem;
to the fool, a solution."
– Marcus Aurelius*

This chapter introduces and investigates a novel class of models; these are attractive models with varying local potentials that we term *patch-potential models*. The inspiration for this chapter is twofold: First, our solution space analysis was restricted to relatively small models so far, which limits the insights; this fueled the search for models that exhibit more than three fixed points. Second, the proof for optimality of SBP relied on the fact that the global minimum of the Bethe free energy provides the most accurate marginals. The reliance on this crucial detail led to questioning the common assumption that such a correspondence extends to more general models as well.

We begin with the introduction of patch potential models and the specification of the considered models in Section 7.2. The main contribution of Section 7.3 is the discussion of the solution space properties by means of an example. In particular, this includes a detailed evaluation of the error in the marginals and the partition function for a wide range of parameters. We then provide theoretical arguments in Section 7.4 that explain the differences between accurate marginals and an accurate partition function and show under which conditions a fixed point exists that approximates both quantities well.

The content of this chapter constitutes the major part of the work in [KP19].

7.1 Motivation

We have restrained ourselves to relatively well-understood models so far: both the exhaustive analysis of the solution space in Chapter 5 and the validation of the proposed algorithm in Chapter 6 are limited to models with identical or random potentials. While models with identical potentials either have one or three fixed points – and are thus too restrictive as to fully understand the behavior of BP – models with random potentials lend themselves to an analysis in terms of expected behavior. The randomness of the potentials, however, introduces a solution space so complex that it cannot be analyzed exhaustively anymore.

We would like to have a model class, simple enough so that it can be well understood; yet complex enough so that it exhibits a rich solution space with more than just three fixed points. For this particular purpose, we introduce *patch potential models*, a rich class of attractive models with inherent structure. These models exhibit many interesting phenomena and provide deep insights into the relationship between the approximation quality of the marginals and the partition function.

So far we have observed that the approximation quality may be severely affected by the existence of multiple fixed points with varying accuracy. In the optimization literature, this issue is often eluded by obtaining and combining all fixed points (cf. RSB theory in Section 3.5.4).

Although a similar approach seems promising for the models studied so far (cf. Chapter 5), computing all fixed points is problematic for more general models. It is often simply not possible to obtain the set of all fixed points in an efficient manner. Remember that one particular aim of Chapter 5 was to specify model classes for which a unique fixed point exists; with the focus on patch potential models, we adhere to this tradition but aim to specify model classes for which a structured and well-behaved solution space is present. Instead of conditions for uniqueness, we are interested in conditions that allow the set of all fixed points to be obtained efficiently.

On the other hand, if one is only interested in selecting a single fixed point, one should obviously strive for selecting the best possible fixed point available. Unfortunately, however, fixed points cannot be compared with respect to the marginal accuracy unless the exact solution is available.

Considering the equivalent variational interpretation admits various provable convergent algorithms (cf. Section 3.3.3). Note that the correspondence to the Bethe free energy allows one to make some more quantitative statements regarding the approximation quality. At least for attractive models, it is well established that the Bethe free energy upper bounds the Gibbs free energy [Ruo12] so that the global minimum provides the most accurate approximation of the partition function. Similar properties are not known for the marginal accuracy, however, and, except for rather simple models (as analyzed in Chapter 5), it remains an open question whether accurate marginals are to be obtained at the global minimum of the Bethe free energy. It is a common conjecture that this is the case; if true, this would provide a simple way of comparing the marginal accuracy of multiple solutions in terms of their corresponding partition function. This chapter studies this relationship in detail and further provides sufficient conditions for the global minimum of $\mathcal{F}_{\mathcal{B}}$ to yield the most accurate marginals.

7.2 Model Specifications

We focus on one specific model class that represents a special case of binary pairwise models. We will consider only finite-size attractive models i.e., where all couplings $J_{ij} > 0$ are positive; specifically, we consider models with equal couplings $J_{ij} = J$ for all edges $(i, j) \in \mathbf{E}$. We shall further distinguish three different types of attractive models that show increasingly complex behavior: (i) attractive models with vanishing local fields $\theta_i = 0$; (ii) attractive models with unidirectional fields, i.e., either $\theta_i < 0$ or $\theta_i > 0$; and (iii) finally, attractive models with arbitrary local fields. Such models are particularly interesting in terms of their phase transitions and are studied under the name of ferromagnetic random-field Ising models (RFIM) in physics where all θ_i are drawn according to some distribution.

7.2.1 Definitions

Attractive models with vanishing fields either have a unique or two *symmetric* fixed points both for infinite-size models [MM09] as well as for finite-size models (cf. Chapter 5). The marginals of two fixed points m and k are considered as *symmetric* if

$$\tilde{P}_{X_i}^m(+1) = 1 - \tilde{P}_{X_i}^k(+1) \quad (7.1)$$

for all X_i . An eminent consequence of how the Bethe free energy was defined as a function of the pseudomarginals (cf. (3.24)) is that symmetric fixed points must also have the same value of $\mathcal{F}_{\mathcal{B}}$. Attractive models with unidirectional fields show a similar behavior and – although not exactly symmetric – have two fixed points that are almost symmetric.

Another important concept are *flipped* random variables: a random variable is flipped if the

marginals are not aligned with the local potential, i.e., if

$$\left(\frac{\tilde{P}_{X_i}(+1)}{\tilde{P}_{X_i}(-1)} - 1 \right) \theta_i < 0. \quad (7.2)$$

We further say that a fixed point is *state-preserving* if no random variable is flipped. If all marginals are in favor of the same state x_i , i.e., if $\tilde{P}_{X_i}(x_i) > 0.5$ for all $X_i \in \mathbf{X}$ we call the corresponding fixed point *biased towards* x_i .

Attractive models with arbitrary local fields exhibit many non-trivial properties, may have a complex solution space, and are studied as one of the simplest forms of disordered systems [You98]. Disordered systems are systems that potentially have many fixed points, whereas many random variables are flipped.

7.2.2 Patch Potential Models

The definition of patch potential models follows the definitions of the RFIM, with the main difference that the local potentials are not i.i.d but obey a correlation between neighboring random variables. Moreover, we will only consider models with identical values for all local fields, albeit possibly with different sign, i.e., $\theta_i \in \{-\theta, +\theta\}$.

Definition 7.2.1. Patch potential models are binary pairwise models in accordance with (2.14) that have attractive couplings $J_{ij} = J > 0$ and that consist of multiple non-overlapping patches \mathcal{G}_i with $\mathcal{G} = \bigcup_i \mathcal{G}_i$. A patch $\mathcal{G}_i = (\mathbf{X}_i, \mathbf{E}_i)$ is a connected subgraph that is induced by a subset of nodes $\mathbf{X}_i \subset \mathbf{X}$ with identical local potentials $\theta_i = \theta$ or $\theta_i = -\theta$, where $\mathbf{E}_i = \{(i, j) \in \mathbf{E} : X_i, X_j \in \mathbf{X}_i\}$.

Note that we will only consider models with sufficiently large patches, so that the exact marginals are state-preserving. Let us first consider a minimal example that is rich enough to exhibit some non-trivial (i.e., non-symmetric) fixed points while being structured enough to admit only few fixed points. This example serves as a model that allows us to get some intuition (cf. Section 7.3) before we discuss the properties of patch potential models in a more general manner (cf. Section 7.4).

Example 10 (Patch Potential Model).

Let $\mathcal{G} = (\mathbf{X}, \mathbf{E})$ be a regular two-dimensional grid graph of size $n \times n$ with two equal-sized patches. All variables in \mathcal{G}_1 experience a positive local field $\theta_1 = \theta$ whereas all variables in \mathcal{G}_2 experience the same negative local field $\theta_2 = -\theta$ (cf. Figure 7.1).




Figure 7.1: Exact solution for Example 10. Nodes are depicted in orange if $P_{X_i}(X_i = 1) > 0.5$ and in blue otherwise; the opacity illustrates the value of the marginals.

The patch potential model is especially appealing as the composition of relatively few patches admits a simplified treatment and comes with a couple of beneficial properties. In particular, we

can identify a region in the parameter space (θ, J) that features a structured and well-behaved solution space (cf. Section 7.3.2).

7.3 Fixed Point Behavior

If BP converges, it often provides accurate results; if multiple fixed points exist, however, the performance may vary considerably between different fixed points. We briefly discuss the RSB (replica symmetry breaking) assumption that expresses the exact marginals as a combination of all fixed points and illustrate why its success is limited to optimization problems so far (Section 7.3.1).

Then, we discuss the solution space of Example 10 over a range of parameters and specify different regions according to the structure of the solution space (Section 7.3.2).

Assessing the approximation quality of a specific fixed point is required to state performance guarantees of BP. We recap existing results (Section 7.3.3) and discuss how the error of the pseudomarginals and the Bethe partition function are related for patch potential models (Section 7.3.4).

7.3.1 Combination Of Fixed Points

(Non-) convexity of the Bethe free energy depends on the structure of the graph and the potentials. If the model has loops and sufficiently strong couplings multiple local fixed points will exist (cf. Chapter 5). The common narrative considers the existence of multiple fixed points as particularly problematic. While this arguably creates a scenario where the performance of BP is more inconsistent, it also opens the door for methods that rely on the RSB assumption (cf. Section 3.5.4). Remember that this allows us to form the exact solution according to

$$P_{X_i}(x_i) = \frac{1}{\sum_m \mathcal{Z}_B^m} \sum_{m=1}^M \mathcal{Z}_B^m \tilde{P}_{X_i}^m(x_i),$$

if all local minima of \mathcal{F}_B are known and $\mathbf{M} = \{(\mathcal{Z}_B^1, \tilde{P}_B^1), \dots, (\mathcal{Z}_B^M, \tilde{P}_B^M)\}$ is available. Note that the RSB-assumption is well-established for infinite-size models [MPV87]; the validity of it for finite-size models, however, remains controversial. Yet, the doubt about finite-size models primarily stems from the lack of rigorous analysis and apart from that, the RSB assumption is empirically well-confirmed for finite-size models as well (an exhaustive overview on recent developments is presented in [LCMRTR13]). Additionally, we computed the exact solution for all considered models using the junction tree algorithm and confirmed the RSB assumption for the patch-potential models as well.

One efficient way to evaluate (3.53) for constrained satisfaction problems is known as survey propagation [BMZ05]. The extension to more general models, however, still remains somewhat elusive.

Approximate Survey Propagation

We have seen in Chapter 5 that a convex combination of the fixed points yields the exact solution in the presence of multiple fixed points. Despite such promising results, the estimation of the marginals according to this combination is hindered by the need for all M solutions. Obtaining all fixed points that correspond to local minima of the Bethe free energy is a complex task only possible for small-scale models (as in Chapter 5) and models with certain structure (e.g., random graphs [COP19]), or potential-type (e.g., for optimization problems [ZK16]).




Figure 7.2: P_μ^m is the fraction of BP runs that converge to fixed point m with the corresponding Bethe partition function \mathcal{Z}_B^m . The mismatch increases with N and θ .

An approximate version of survey propagation was recently applied to similar models as in this work [SRF16]. This was achieved by assuming that the fraction of randomly initialized BP runs P_μ^m that converges to the m^{th} fixed point provides an approximation of the partition function \mathcal{Z}_B^m . This assumption is valid for attractive models with vanishing local fields; yet it is unclear how this generalizes to models with non-vanishing local fields.

We aim to validate the assumption for regular grid graphs with $n \times n$ variables, $\theta \neq 0$, and with couplings large enough to admit two fixed points. Therefore, we compare both measures for both fixed points by relating the ratio between the partition functions $\mathcal{Z}_B^1/\mathcal{Z}_B^2$ to the ratio P_μ^1/P_μ^2 . The log-ratio⁴⁶ between both measures is depicted in Figure 7.2. One would expect a constant value close to zero if P_μ^m provides a good estimate of \mathcal{Z}_B^m ; this is obviously not the case as $\mathcal{Z}_B^1/\mathcal{Z}_B^2$ grows more rapidly. We conclude that the fraction of convergent BP runs serves as a poor estimate of the partition function with the consequence that an approximate evaluation of (3.53) leads to inaccurate marginals. This is particularly true as the local field and the model size increase.

This raises two immediate questions: (i) Can we specify certain model-structures or parameter configurations that grant efficient methods to obtain all fixed points in order to evaluate (3.53)? (ii) If we obtain a subset of all fixed points $\hat{\mathbf{M}} \subset \mathbf{M}$, can we compare the available fixed points and select the best one?

7.3.2 Solution Space of Patch Potential Models

The solution space for a wide range of patch potential models is analyzed to answer whether parameter configurations exist for which all fixed points can be obtained efficiently. A more formal analysis that explains the subsequent observations is presented in Section 7.4.

Let \mathcal{G} be a 10×10 grid graph with two equal-sized patches (Example 10). This model exhibits three different regions, separated by critical values $J_A(\theta)$ and $J_C(\theta)$; see Figure 7.3 and Figure 7.4 for an illustration of the decomposition into multiple fixed points according to (3.53).

A unique fixed point exists for $J < J_A(\theta)$, i.e., inside region (I), and BP converges; this fixed point is state-preserving but slightly overestimates the marginals (cf. Section 7.3.4). Additional fixed points emerge inside region (II) as the coupling strength increases to $J_A(\theta) < J < J_C(\theta)$. There are three fixed points (cf. Theorem 12) and all three fixed points are stable (cf. Theorem 13). These fixed points consist of two symmetric fixed points where all marginals favor one

⁴⁶ The log-ratio is independent of the coupling strength as long as J is large enough to admit two fixed points.




Figure 7.3: Illustration of the fixed points for all regions. The circle-width corresponds to the value of \mathcal{F}_B .

particular state and one state-preserving fixed point (cf. Section 7.3.4). As the coupling strength increases even further to $J > J_C(\theta)$, i.e., inside region (III), all three fixed points remain but are suddenly accompanied by many more fixed points. It will therefore be increasingly hard to obtain all fixed points numerically, so that one can only hope to obtain a subset of all fixed points in practice.

The actual boundaries between the regions are numerically estimated and are depicted in Figure 7.4. The fixed points are obtained by repeated application of BP (2000 times for each (θ, J)) with different random initial conditions. Furthermore, we apply random scheduling to enhance the convergence properties as any predetermined schedule would favor a specific fixed point.

To answer question (i) from Section 7.3.1: one region exists in the parameter space (illustrated in blue) for which all fixed points can be obtained efficiently. For region (III) (illustrated in red), however, the number of fixed points suddenly increases and we cannot rely on BP to obtain all fixed points.

7.3.3 Approximation Accuracy

Let $J > J_c(\theta)$ and assume that a subset $\tilde{\mathbf{M}} \subset \mathbf{M}$ of all fixed points that constitute minima of \mathcal{F}_B is provided; then, how can we select the best one? Unfortunately, there is no way to tell us how accurate a particular fixed point is (if we do not have access to the exact solution). It is therefore an important problem in its own to measure the accuracy, or at least provide a bound on the approximation error. We will first discuss established results regarding the accuracy of both the Bethe partition function and the pseudomarginals. Subsequently, we will delve into the particularities for patch potential models and show how the accuracy may differ between both objectives.

We measure the error of the partition function $\mathcal{Z}_B^m = \mathcal{Z}_B(\tilde{P}_B^m)$ of the m^{th} fixed point by $E_{\mathcal{Z}}(m)$ in terms of the relative error between the log-partition functions according to (3.36). Note that we only consider attractive models, for which the Bethe partition function also bounds the partition function, i.e., $\mathcal{Z}_B < \mathcal{Z}$ [Ruo12]; obtaining the global minimum of \mathcal{F}_B is therefore optimal with respect to the error of the partition function as $\arg \min_{\mathcal{Z}_B^m} (E_{\mathcal{Z}}(m)) = \exp(-\min_{\mathbf{L}} \mathcal{F}_B(\tilde{P}_B^m))$.

The error of the singleton marginals $E_P(m)$ is measured by the mean squared error (MSE) according to (3.38). We are not aware of an explicit relationship that connects both worlds and relates the error of the marginals to the error of the partition function, except for homogeneous⁴⁷ attractive models (cf. Lemma 5). It is therefore often assumed that minimizing \mathcal{F}_B will be optimal in terms of marginal accuracy for more general models as well (cf. Chapter 6

⁴⁷ These are models that have a single value J for all edges and a single value θ for all variables




Figure 7.4: Illustration of all regions and boundaries for Example 10: the black dots depict the boundary below which (7.3) holds; the approximated boundary according to (7.4) is depicted by the solid black line.

and [WTJS14]). More formally this assumption states that

$$\arg \min_{m \in \mathbf{M}} E_{\mathcal{Z}}(m) = \arg \min_{m \in \mathbf{M}} E_P(m). \quad (7.3)$$

This assumption is, however, not valid in general as we will now show in Section 7.3.4.

7.3.4 Comparison Of Marginal Accuracy and Partition Function Accuracy

We aim to evaluate the relationship between the accuracy of the pseudomarginals and the accuracy of the partition function and whether (7.3) holds in general. First, we state that (7.3) does hold for homogeneous attractive models that have two fixed points at most [WTJS14]. This is a direct consequence of (3.53) or by casting the model as a larger one with vanishing local fields (cf. discussion after Theorem 9).

Lemma 5. *Attractive models with identical values $\theta_i = \theta$ have two fixed points for $J > J_A(\theta)$. The fixed point m that minimizes $E_{\mathcal{Z}}(m)$ further provides the global minimum $\min_{\perp}(\mathcal{F}_{\mathcal{B}})$ and minimizes $E_P(m)$ as well.*

Second, we empirically validate whether minimizing $\mathcal{F}_{\mathcal{B}}$ will provide the most accurate marginals for Example 10. Figure 7.5 illustrates the error in the marginals and the error in the partition function for all fixed points. The fixed point that provides the global minimum to $\mathcal{F}_{\mathcal{B}}$, and thus minimizes $E_{\mathcal{Z}}(m)$ is emphasized in blue, the fixed point minimizing $E_P(m)$ is emphasized in red, whereas the fixed point minimizing both quantities jointly is emphasized in green.

Let us take a closer look at region (II) in particular: three fixed points exist that can be combined to yield the exact solution (see Figure 7.1 for the exact solution). Two of these fixed points, r and q , are each biased towards one state and, because of the symmetric model, have identical values $\mathcal{F}_{\mathcal{B}}^r = \mathcal{F}_{\mathcal{B}}^q$. The state preserving fixed point p on the other hand provides the most accurate marginals inside (II). However, while p also provides the global minimum of $\mathcal{F}_{\mathcal{B}}$ for small values of J , Figure 7.5 shows that $\mathcal{F}_{\mathcal{B}}^p$ turns into a local minimum for $J \geq 0.65$. No principle relationship between the accuracy of the marginals and the partition function can therefore be observed inside (II) and (7.3) does not necessarily hold (cf. Theorem 15).

For region (III) many more fixed points (u, v, \dots) emerge that all have similar values $E_{\mathcal{Z}}(u)$ and $E_P(u)$; we visualize some of them in Figure 7.5. These fixed points provide slightly more




Figure 7.5: Accuracy of the marginals (a) and of the partition function (b) for Example 10 with $|\theta_i| = 0.1$: we emphasize the fixed points minimizing $E_Z(m)$ (blue), minimizing $E_P(m)$ (red), and minimizing both quantities (green).

accurate marginals than the state-preserving one, although it should be noted that all fixed points do not approximate the marginals well inside (III). On the contrary, considering $E_Z(u)$, these additional fixed points provide the worst approximation to the partition function and have even higher values $\mathcal{F}_B^u > \mathcal{F}_B^p > \mathcal{F}_B^q$. The biased fixed points p, q that approximate the marginals worst, on the other hand, approximate the partition function relatively well.

Why fixed points exist that minimize the marginal error but are only local minima of \mathcal{F}_B can, however, not be answered by the above observations. Closer inspection of \mathcal{F}_B for different types of fixed points reveals a threshold (black dots in Figure 7.4) below which (7.3) holds. Some mild assumptions on the solution space lead to a lower bound on this threshold (cf. Theorem 16) according to

$$2J\sqrt{N} - \theta N = 0. \quad (7.4)$$

This bound, illustrated by the solid black line in Figure 7.4, becomes asymptotically exact. Note that the slope, defined by (7.4) increases with the model size N so that the global minimum of \mathcal{F}_B provides the most accurate marginals for a wider range of parameters.

7.4 Theoretical Analysis

Here we properly define the boundaries $J_A(\theta)$ and $J_C(\theta)$ between different regions and provide formal arguments that explain the observations from Section 7.3.2. While some properties are directly attributable to (3.53), several results are based on the fact that the patch potential model consists of multiple patches with a unidirectional local field. First, we need to prepare an alternative update equation that makes the interactions between two patches more explicit. For that purpose, we will introduce an effective field that acts on the boundary of each patch and incorporates the influence from all other patches.

We only present the most insightful proofs below and defer some longer proofs to the Appendix A.2. Additionally, we prepare some corollaries that simplify the results for models with two equal-sized patches as in Example 10.

7.4.1 Effective Field

We introduce an effective field $\tilde{\theta}_i$ for all variables that lie on the patch-boundary to incorporate the interactions with the neighboring patches.

Theorem 11 (Effective Field). *Let X_i be a variable on the boundary of patch \mathbf{X}_i that receives messages from inside, i.e., $X_k \in \mathbf{X}_i$, and outside, i.e., $X_j \in \{\mathbf{X} \setminus \mathbf{X}_i\}$, the patch. The effective field $\tilde{\theta}_i$ acts on the boundary according to*

$$\tilde{\theta}_i = \theta_i + \sum_{X_j \in \{\partial(i) \setminus \mathbf{X}_i\}} \operatorname{atanh}(2\mu_{ji}(X_i = 1) - 1). \quad (7.5)$$

Messages from outside the patch are now subsumed by $\tilde{\theta}$ and the additive terms in (7.5) will be positive if $\mu_{ji}(X_i = 1) > \mu_{ji}(X_i = 0)$ and negative otherwise. This is particularly important in the definition of the region boundaries and admits “independent” treatment of every patch.

7.4.2 Definition of Region (II)

The notion of an effective field (Theorem 11) allows us to define the boundaries between the three distinct performance regions of patch potential models. We discuss the solution space in detail and what can be said about the performance of BP. Let us denote the second region, i.e., the region where the global behavior can be inferred by treating the patches individually by $(II) = \{\theta, J\}$.

Definition 7.4.1 (Region). A parameter set $(\theta, J) \in (II)$ if and only if the following conditions are satisfied:

- (1.) Let $J_A(\mathcal{G}_i, \theta)$ denote the critical value for the couplings beyond which multiple fixed points exist.⁴⁸ Then every patch $\mathcal{G}_i \in \mathcal{G}$ must have its respective threshold below the actual coupling strength, i.e., $J_A(\mathcal{G}_i, \theta) < J$
- (2.) Consider all pairs of patches \mathcal{G}_i and \mathcal{G}_j ; if one patch, e.g., \mathcal{G}_i has its variables flipped, the imposed effective field on the boundary must stabilize the second patch \mathcal{G}_j so that $J < J_A(\mathcal{G}_j, \tilde{\theta}) = J_C(\mathcal{G}_j, \theta)$.

These conditions implicitly define the “well-behaved” region (II). Definition 7.4.1.1 provides the lower boundary of region (II) as only a unique fixed point would exist otherwise. It may be less obvious how Definition 7.4.1.2 provides the upper boundary of region (II). Note that $J < J_A(\mathcal{G}_j, \tilde{\theta})$ is a necessary condition if \mathcal{G}_i is flipped, as parts of \mathcal{G}_j would flip otherwise and lead to disordered behavior (cf. Figure 7.5). The restriction to (II) and the exclusion of disordered solutions further validates the RSB assumption [MM09, Chapter 19].

7.4.3 Properties Of Region (II)

In this chapter, we are particularly interested in understanding the properties of BP inside region (II) that complies with the subsequent properties. Note that the properties inside region (II) are a direct consequence of Definition 7.4.1. Many arguments will rely on the fact that every attractive model with vanishing or unidirectional local fields either has a unique stable or two stable fixed points (cf. Section 7.2).

We will first bound the number of possible BP fixed points.

Theorem 12 (Existence). *Let \mathcal{U} be a patch potential model with $(\theta, J) \in (II)$. The amount of fixed points M grows with the number of patches (rather than the number of variables). Specifically, we have $M = \mathcal{O}(2^{(|\mathcal{G}_i|)})$, where $|\mathcal{G}_i|$ denotes the number of patches.*

⁴⁸ Note that an analytical solution only exists for graphs with vanishing fields of infinite size or periodic boundary conditions, but the threshold can be estimated numerically.

Proof. First, assume that a given patch \mathcal{G}_i is flipped; then by the definition of the patch potential model (Definition 7.2.1) and by Theorem 11 it follows that effective field $\tilde{\theta}$ is aligned with the local field of the variables at any neighbor patch \mathcal{G}_j so that

$$\text{sgn}(\tilde{\theta}_j) = \text{sgn}(\theta_j), \quad (7.6)$$

$$|\tilde{\theta}_j| > |\theta_j|. \quad (7.7)$$

Further, let us recall the definition of (II) (Definition 7.4.1.2 in particular). It follows that the effective field stabilizes its neighbor patch \mathcal{G}_j , i.e., $J_A(\mathcal{G}_j, \tilde{\theta}) > J_A(\mathcal{G}_j, \theta)$ so that, according to the definition $J < J_A(\mathcal{G}_j, \tilde{\theta})$, and \mathcal{G}_j admits only a unique solution.

Second, assume that \mathcal{G}_i is not flipped; then it follows by Theorem 11 that $\tilde{\theta}_i < \theta_i$ for X_i on the boundary of the neighbor patch \mathcal{G}_j . This decrease in the local field reduces the threshold for the existence of two solutions to smaller values of J (cf. [KP17]) so that

$$J_A(\mathcal{G}_j, \tilde{\theta}) < J_A(\mathcal{G}, \theta) < J. \quad (7.8)$$

By Definition 7.4.1.1 it follows that the neighbor patch \mathcal{G}_j has two fixed points now, and it depends on the initialization to which one BP will converge.

Finally, we aim to show that the number of possible fixed points is bounded. Therefore we want to stress that the above arguments show how patches can either be aligned with the local potential or be flipped; it is crucial that every patch acts as one instance and that all variables belonging to one patch are aligned. Else the fixed point would be disordered which we rule out precisely by Definition 7.4.1. This and the fact that we are considering binary random variables limits the number of possible solutions to

$$M \leq 2^{|\mathcal{G}_i|}, \quad (7.9)$$

where $|\mathcal{G}_i|$ denotes the overall number of patches. □

Corollary 12.1 (Example 10). *Let \mathcal{U} be a patch potential models with two equal-sized patches (cf. Example 10). Then, for $(\theta, J) \in (II)$ three fixed points exist; these are one state preserving fixed point and two fixed points that have all variables biased towards one of both states. Note that both patches can not be flipped simultaneously inside (II) (cf. proof of Theorem 12) as one patch would stabilize, i.e., prohibit from flipping, the second patch.*

Theorem 12 is of great practical relevance for the RSB assumption (3.53), i.e., whether a combination of BP fixed points can form the exact solution. The fact that there is a relatively small number of fixed points makes the task of obtaining them practically feasible. Existence alone, however, is not sufficient as we have to rely on some numerical method that obtains all fixed points; if we aim to apply BP for that matter there is the additional requirement for all fixed points to be stable. Fortunately, it turns out that all fixed points inside (II) are stable indeed.

Theorem 13 (Stability). *Let \mathcal{U} be a patch potential model with $(\theta, J) \in (II)$. Then, every fixed point \tilde{P}_B^m is a stable fixed point for BP.*

Finally, as an immediate consequence of the limited amount of fixed points (Theorem 12), all of which are stable (Theorem 13), it follows that the exact solution can be computed according to (3.53) in practice. One can for example apply BP repeatedly, possibly in parallel, with random initialization to obtain and combine all fixed points.

Marginal Accuracy

Theorem 14 (Marginal Accuracy). *The MSE of the singleton marginals $E_P(k)$ of the k^{th} solution \tilde{P}_B^k relates to the ratio of the Bethe partition functions according to*

$$E_P(k) = \frac{2}{N(\sum_m \mathcal{Z}_B^m)^2} \sum_{X_i \in \mathbf{X}} \left| \sum_{m \setminus k} \mathcal{Z}_B^m (\tilde{P}_{X_i}^m - \tilde{P}_{X_i}^k) \right|^2.$$

Proof. According to (3.53) we can express the exact solution by the convex combination of all fixed points. Consequently, using symmetry properties of the binary random variables, the error is given by

$$\begin{aligned} E_P(k) &= \frac{2}{N} \sum_{X_i \in \mathbf{X}} \left| \frac{\sum_m \mathcal{Z}_B^m \tilde{P}_{X_i}^m}{\sum_m \mathcal{Z}_B^m} - \tilde{P}_{X_i}^k \right|^2 \\ &= \frac{2}{N} \sum_{X_i \in \mathbf{X}} \left| \frac{\sum_m \mathcal{Z}_B^m \tilde{P}_{X_i}^m - \sum_m \mathcal{Z}_B^m \tilde{P}_{X_i}^k}{\sum_m \mathcal{Z}_B^m} \right|^2 \\ &= \frac{2}{N(\sum_m \mathcal{Z}_B^m)^2} \sum_{X_i \in \mathbf{X}} \left| \sum_{m \setminus k} \mathcal{Z}_B^m \tilde{P}_{X_i}^m - \tilde{P}_{X_i}^k \right|^2, \end{aligned} \quad (7.10)$$

where we first, bring everything on the same denominator so that the k^{th} contribution cancels out subsequently. \square

Representing the MSE according to Theorem 14 is particularly appealing as it omits the need for expressing the exact marginals. This further provides a way to express the ratio of the marginal error between two fixed points.

Corollary 14.1. *The MSE-ratio of two fixed points k and l is a ratio of weighted partition functions according to*

$$\frac{E_P(k)}{E_P(l)} = \frac{\sum_{X_i \in \mathbf{X}} \left| \sum_{m \setminus k} \mathcal{Z}_B^m (\tilde{P}_{X_i}^m - \tilde{P}_{X_i}^k) \right|^2}{\sum_{X_i \in \mathbf{X}} \left| \sum_{m \setminus l} \mathcal{Z}_B^m (\tilde{P}_{X_i}^m - \tilde{P}_{X_i}^l) \right|^2}. \quad (7.11)$$

Expressing the ratio of the marginal error according to (7.11) is advantageous in elaborating on the difference between accuracy of the approximated marginals and the approximated partition function. We define the mismatch between $\tilde{P}_{X_i}^m$ at two fixed points k and l by

$$Q_i(k, l) = \tilde{P}_{X_i}^k(X_i = 1) - \tilde{P}_{X_i}^l(X_i = 1). \quad (7.12)$$

Now, let us denote the error of the state preserving fixed point by $E_P(p)$ and of the fixed point that has all marginals biases towards $x_i = 1$ by $E_P(q)$. Then – maybe non-surprising as the exact solution is state preserving as well – we show that the state-preserving fixed point has the most accurate marginals. We first discuss the error-ratio in a general manner in Theorem 15 before considering the special case of a model with two patches, i.e., for Example 10.

Theorem 15 (Error Ratio). *Let \mathcal{U} be a patch potential model with $(\theta, J) \in (II)$. The state preserving fixed point p provides more accurate marginals than the fixed point q that has all marginals biased to one state, i.e.,*

$$\frac{E_P(p)}{E_P(q)} < 1. \quad (7.13)$$

Proof. To show that, irrespective of the value of \mathcal{Z}_B , $\frac{E_P(p)}{E_P(q)} < 1$ we assume that the state preserving fixed point does not minimize the Bethe free energy, i.e., $\mathcal{Z}_B^q > \mathcal{Z}_B^p$.

Without loss of generality, we make some prior assumptions on the model: First, we assume that the overall number of variables with a positive local field equals the number of variables with a negative local field, i.e.,

$$|\{X_i \in \mathbf{X} : \theta_i = +\theta\}| = |\{X_j \in \mathbf{X} : \theta_j = -\theta\}|. \quad (7.14)$$

Second, we assume that all patches are of equal size. And finally, we group all possible fixed points according to their marginals and denote them as follows: the state-preserving fixed point is referred to as p ; all fixed points that have more patches biased towards $X_i = +1$ are referred to as $s = 1, \dots, S$, with q being the fixed point that has all variables biased towards $X_i = +1$; all fixed points that have more patches biased towards $X_i = -1$ are referred to as $t = 1, \dots, T$, with r being the fixed point that has all variables biased towards $X_i = -1$.

This has some implications that will ease the subsequent analysis significantly. Specifically, the number of fixed points favoring one state equals the number of fixed points favoring the other state, i.e., $S = T$ by (7.14) and by Theorem 12.

Another important consequence of incorporating the interactions between patches into an effective field is that the number of variables that favor one state has an immediate influence on the value of the singleton marginals. It can be shown that the effective field, if stronger than the local field – note that the effective field is stronger than the local field whenever two neighboring patches are biased towards the same state – increases the bias of the variables. This intuitive statement is a consequence of the Griffiths-Hurst-Sherman inequality [Gri67] that can be extended to specific fixed points by straightforward manipulations (cf. [KKP18]). In essence this means that for all variables $X_i \in \mathbf{X}$ we have

$$\tilde{\mathcal{Z}}_{X_i}^q(X_i = 1) \geq \tilde{\mathcal{P}}_{X_i}^s(X_i = 1) \geq \tilde{\mathcal{P}}_{X_i}^p(X_i = 1) \geq \tilde{\mathcal{P}}_{X_i}^t(X_i = 1) \geq \tilde{\mathcal{P}}_{X_i}^r(X_i = 1). \quad (7.15)$$

Moreover, as all patches have equal size and because $J_{ij} = J$ as well as $\theta_i \in \{-\theta, \theta\}$ every fixed point s has a symmetric fixed point t that has the same value for the approximate partition function (cf. Proof of Theorem 16). That is, except for the state-preserving fixed point p all fixed points come in couples that satisfy

$$\mathcal{Z}_B^s = \mathcal{Z}_B^t. \quad (7.16)$$

We will further utilize the properties of the mismatch $Q_i(k, l)$; in particular the symmetry property

$$Q_i(k, l) = -Q_i(l, k), \quad (7.17)$$

and the expansion property

$$Q_i(k, l) = Q_i(k, m) + Q_i(m, l). \quad (7.18)$$

Finally, we express the error ration between the state-preserving fixed point p and the biased

fixed point q according to Corollary 14.1 so that

$$\begin{aligned}
\frac{E_P(p)}{E_P(q)} &= \frac{\sum_{X_i \in \mathbf{X}} \left| \sum_{m \setminus p} \mathcal{Z}_B^m Q_i(m, p) \right|^2}{\sum_{X_i \in \mathbf{X}} \left| \sum_{m \setminus q} \mathcal{Z}_B^m Q_i(m, q) \right|^2} \\
&\stackrel{(a)}{=} \frac{\sum_{X_i \in \mathbf{X}} \left| \sum_s \mathcal{Z}_B^s Q_i(s, p) + \sum_t \mathcal{Z}_B^t Q_i(t, p) \right|^2}{\sum_{X_i \in \mathbf{X}} \left| \sum_{s \setminus q} \mathcal{Z}_B^s Q_i(s, q) + \sum_t \mathcal{Z}_B^t Q_i(t, q) + \mathcal{Z}_B^p Q_i(p, q) \right|^2} \\
&\stackrel{(b)}{=} \frac{\sum_{X_i \in \mathbf{X}} \left| \sum_{s \setminus q} \mathcal{Z}_B^s (Q_i(s, p) + Q_i(t, p)) + \mathcal{Z}_B^q (Q_i(q, p) + Q_i(r, p)) \right|^2}{\sum_{X_i \in \mathbf{X}} \left| \sum_{s \setminus q} \mathcal{Z}_B^s (Q_i(s, q) + Q_i(t, q)) + \mathcal{Z}_B^q Q_i(r, q) + \mathcal{Z}_B^p Q_i(p, q) \right|^2}, \tag{7.19}
\end{aligned}$$

where (a) follows from splitting the sum into the fixed points s that are more biased towards $X_i = 1$ and into the fixed points t that are more biased towards $X_i = -1$. Note that the state-preserving fixed point p does not belong to either set and is consequently expressed explicitly in the denominator. For (b) we make use of (7.16) and arrange the terms by making the dependence on q and r explicit so that the sum goes over the same terms in the numerator and in the denominator.

We can further express the error ratio and bound it using Jensen's inequality according to

$$\begin{aligned}
\frac{E_P(p)}{E_P(q)} &\stackrel{(a)}{=} \frac{\sum_{X_i \in \mathbf{X}} \left| \sum_{s \setminus q} \mathcal{Z}_B^s (Q_i(s, p) + Q_i(t, p)) + \mathcal{Z}_B^q (Q_i(q, p) + Q_i(r, p)) \right|^2}{\sum_{X_i \in \mathbf{X}} \left(\sum_{s \setminus q} \mathcal{Z}_B^s |Q_i(s, q) + Q_i(t, q)| + \mathcal{Z}_B^q |Q_i(r, q)| + \mathcal{Z}_B^p |Q_i(p, q)| \right)^2} \\
&\leq \frac{\sum_{X_i \in \mathbf{X}} \left(\sum_{s \setminus q} \mathcal{Z}_B^s |Q_i(s, p) + Q_i(t, p)| + \mathcal{Z}_B^q |Q_i(q, p) + Q_i(r, p)| \right)^2}{\sum_{X_i \in \mathbf{X}} \left(\sum_{s \setminus q} \mathcal{Z}_B^s |Q_i(s, q) + Q_i(t, q)| + \mathcal{Z}_B^q |Q_i(r, q)| + \mathcal{Z}_B^p |Q_i(p, q)| \right)^2}, \tag{7.20}
\end{aligned}$$

where separating the norm does not change the result in (a) because of (7.15), which implies $Q_i(m, q) < 0$ for all fixed points $m \neq q$.

For completing the proof we make use of the symmetry property (7.17) and the expansion property (7.18) in order to rearrange the terms; in particular note that $Q_i(q, s) + Q_i(q, t) = Q_i(q, s) + Q_i(q, s) + Q_i(s, p) + Q_i(p, t)$, and that $Q_i(q, r) = Q_i(q, p) + Q_i(p, r)$ so that

$$\begin{aligned}
\frac{E_P(p)}{E_P(q)} &\leq \frac{\sum_{X_i \in \mathbf{X}} \left(\sum_{s \setminus q} \mathcal{Z}_B^s |Q_i(s, p) - Q_i(p, t)| + \mathcal{Z}_B^q |Q_i(q, p) - Q_i(p, r)| \right)^2}{\sum_{X_i \in \mathbf{X}} \left(\sum_{s \setminus q} \mathcal{Z}_B^s |Q_i(s, p) + Q_i(p, t) + Q_i(q, s) + Q_i(q, s)| + \mathcal{Z}_B^q |Q_i(q, p) + Q_i(p, r)| + \mathcal{Z}_B^p |Q_i(q, p)| \right)^2}.
\end{aligned}$$

Note that we have applied (7.17) so that every mismatch-term is strictly positive. It is thus straightforward to see, by comparing all terms, that the numerator is strictly smaller than the denominator for every variable $X_i \in \mathbf{X}$ so that

$$\frac{E_P(p)}{E_P(q)} < 1. \tag{7.21}$$

□

In particular for models with two equal-sized patches, we can simplify the error ratio (7.11) considerably.

Corollary 15.1 (Example 10). *Let $d = Q_i(q, r) > 0$, then*

$$\frac{E_P(p)}{E_P(q)} < \frac{\sum_{X_i \in \mathbf{X}} |Z_B^q d|^2}{\sum_{X_i \in \mathbf{X}} |Z_B^q d + Z_B^p Q_i(p, q)|^2} < 1. \quad (7.22)$$

Note that Corollary 15.1 is an immediate consequence of Theorem 15. Additionally we present an alternative proof that admits some intuitive arguments in Appendix A.2.2.

It follows that the state preserving fixed point p minimizes the marginal error inside (II) irrespective of \mathcal{F}_B^p . This has drastic implications and forbids any relationship between the fixed point minimizing the marginal error and the one minimizing the partition function error.

Free Energy Minimizing Fixed Point

However, despite Theorem 15, the question remains where the difference between $E_Z(m)$ and $E_P(m)$ stems from?

We will now answer this question and provide conditions for $\arg \min E_Z(m) = \arg \min E_P(m)$ to be valid. We further present an approximate condition for the state-preserving fixed point p to simultaneously provide the most accurate marginals and minimize \mathcal{F}_B . Therefore, let us define the following variables (cf. Section A.2.3 in the appendix for a formal introduction): \mathbf{E}_P is the set of all boundary edges; \mathbf{E}_C is the set of edges between variables that favor different states; N_f and N_c are the numbers of flipped and non-flipped variables; and ΔS_B is the difference in the entropy between two fixed points.

Theorem 16. *Let us consider the state-preserving fixed point p with \mathcal{F}_B^p and some other fixed point with \mathcal{F}_B^m . Then, $\mathcal{F}_B^p < \mathcal{F}_B^m$ is the global minimum if*

$$2J(|\mathbf{E}_P| - |\mathbf{E}_C|) < \theta(N - N_c + N_f) + \Delta S_B. \quad (7.23)$$

For models with two equal-sized patches and couplings strong enough for the entropy-term to vanish we can further simplify (7.23) significantly and state that:

Corollary 16.1 (Example 10). *The state-preserving fixed point provides the most accurate marginals and the global minimum \mathcal{F}_B^p if $(\theta, J) \in (II)$ and if*

$$2\sqrt{N}J < N\theta. \quad (7.24)$$

Proof. For the specific case of a grid graph with two equal-sized patches (Example 10) we can further simplify the condition from Theorem 16. Therefore, note that for $\mathcal{F}_B^p < \mathcal{F}_B^q$ to be satisfied, we have $N_c = N_f$ and $|\mathbf{E}_P| = \sqrt{N}$ so that (A.28) reduces to

$$2J\sqrt{N} \leq \theta N + \Delta S_B. \quad (7.25)$$

The definition of (II) requires strong interactions J so that the entropy terms in the free energy is small. We can consequently approximate (7.25) by neglecting the entropy terms.

$$2J\sqrt{N} \leq \theta N. \quad (7.26)$$

□

These sufficient conditions for (7.3) provide a guideline when it would be safe to select the fixed point according to the partition function value. This correspondence tends to hold for models with strong local potentials θ and with increased model-size N as shown in Corollary 16.1.

7.5 Conclusion

In Section 6.4 we had relied on the assumption that finding the global minimum of \mathcal{F}_B gives us the most accurate marginals and partition function. This is a strong assumption that – although never formally verified – was generally believed to be true.

In this chapter we aimed to verify whether such a relationship holds in general. Therefore, we studied the relationship between the accuracy of the marginals and the partition function at stationary points of the Bethe free energy, with far-reaching implications for approximate inference methods that operate on the Bethe free energy. To do so, we introduced a new class of models first. These so-called patch potential models exhibit lots of structure that enabled us to analyze the solution space thoroughly. Yet, patch potential models proved to be flexible enough to produce some surprising insights.

First, we elaborated on the existence of a well-behaved region in parameter space for which the number of fixed points depends only on the number of patches (instead of the variables).

Second, this well-behaved region allowed us to assess the correspondence between the accuracy of the marginals and the partition function; we showed that this common assumption is *not* true in general and explained why – instead of the global minimum – the most accurate marginals may be found at a local minimum of \mathcal{F}_B .

Finally, we inspected under which conditions this mismatch arises and further introduced guarantees for the global minimum of \mathcal{F}_B to remain optimal with respect to marginal accuracy. Note that the mismatch is effectively attributable to finite-size effects which might explain why this behavior is not known in the physics literature, where the focus lies on infinite-size – or at least very large – models. Yet, models of relatively small size do play an important role in many applications of BP.

The analysis of patch potential models may serve as a foundation that enables the extension of survey propagation to problems beyond constrained satisfaction problems. So far this has been limited by the potentially huge amount of fixed point. Our definition of the well-behaved region and the study of patch potential models as a whole suggests that it may be very much possible to obtain the set of all fixed points efficiently, as long as the model is sufficiently structured.

Moreover, our newly developed understanding of the relationship between the accuracy of the marginals and the partition function has a notable impact on our theoretical analysis of SBP in Section 6.4. There we have limited our attention to attractive homogeneous models where the accuracy of both quantities can be used in an interchangeable way. Whether it is the minimum (i.e., the Bethe free energy) or the minimizer (i.e., the pseudomarginals) that is responsible for the optimality of SBP needs to be reassessed. It is not clear if it is really the global minimum of \mathcal{F}_B that originates from the start point. On the contrary, our minimal patch potential model of Example 10 shows that the state-preserving fixed point defines the solution path, so that SBP may very well be in favor of obtaining the fixed point with the most accurate marginals. This could be an explanation for the surprisingly accurate results of SBP for general models and would immensely strengthen the practical relevance of SBP as no method is known so far that deliberately aims to obtain fixed point with the most accurate marginals.

8

Discussion and Open Questions

*"There must be some way out of here,
Said the joker to the thief."
– Bob Dylan*

In this thesis, we performed a systematic analysis of BP's solution space. Finding all fixed points of BP is a relevant but hard problem, except for a few restricted models. We took inspiration from a set of diverse scientific fields to tackle this problem: in particular, we framed BP as dynamical systems and formulated the fixed point equations. Drawing from computational mathematics, we established a way to solve the set of fixed point equations and computed all fixed points for a range of problems.

The knowledge of the full solution space then served as a cornerstone for evaluating the accuracy and the convergence properties of BP. Moreover, we related those performance properties to the parameterization of a given model (i.e., the model size and the specifications of the potentials) and explained how the performance of BP changes if we tune the model.

A detailed summary of our main contributions is presented in each chapter's conclusion. Here we summarize our most interesting findings and discuss how we advanced the understanding of BP. Our findings satisfy a twofold interest: they advance the theoretical understanding and suggest various possibilities of enhancing BP's performance.

- From a theoretical perspective, we made the influence of the model parameters on the performance of BP explicit. We sum up our key-insights as follows:

The influence of the model size is hard to pinpoint down. On the one hand, larger models tend to have a higher average degree, which affects the convergence properties slightly negatively. On the other hand, we established a positive effect of the model size on the marginal accuracy.

The influence of the potentials is much more apparent. Strong pairwise potentials influence the performance of BP detrimentally, whereas strong local potentials increase the accuracy and the convergence properties.

Additionally, we demonstrated a considerable accuracy gap between the best possible and the worst possible fixed point of BP. This poses a fundamental problem when deriving performance guarantees as is not obvious to which fixed point BP will converge.

- From a practical perspective, we justified the exploration of multiple fixed points, intending to select one with good approximation quality. We discussed the underlying mechanisms and limitations of established modifications, for example, damping or scheduling.

We further proposed one particular modification that aims to obtain the global minimum of the Bethe free energy. This method, SBP, consistently obtains more accurate marginals than BP and provides accurate marginals for models where BP fails to converge.

Even the minimization of the Bethe free energy, however, has its limitations if one is only interested in the marginals; we exemplified and explained why the most accurate marginals may not be found at the global minimum of the Bethe free energy.

The analysis of BP's solution space addressed some long-standing questions but we are still far away from understanding every aspect of BP. Below we summarize some intriguing questions that emerged while working on this thesis and that were left unanswered. We distinguish three different directions that appear to be worth pursuing in particular:

1. Structure of the Solution Space:

Although this thesis obtained the full solution space for some models and, in doing so, provided novel insights into the properties of BP, computing the solution space for more complex models remains elusive. As we have seen in Chapter 7, however, it is precisely the consideration of more complex models that can lead to fascinating insights.

The extension of the NPHC method to larger models was primarily hindered by the involved computation of the root count. So far, we only relied on some meta-heuristics to compute the root count. Most relevant models, be it grid graphs or models arising in the context of error-correcting codes, however, exhibit a special graph structure that directly determines the structure of the fixed point equations. A first step towards computing the root count for complex models could be to exploit the graph structure and thus reduce the overall complexity of this crucial step. Alternatively, it would be relevant (from a mathematical perspective) to reformulate the problem to an equivalent system of equations with fewer complex solutions and a lower BKK bound. Such problem-tailored approaches would extend the applicability of the NPHC method and, when applied to large models, would have the potential to provide even more insights into the behavior of BP.

Traditionally, one is often interested in specifying regions with a unique fixed point. We have shown that this seems to be a rather limiting point of view as combinations of multiple fixed points often provide strikingly accurate results. Although obtaining the set of all fixed points is, in general, a hard problem, we have seen in Chapter 7 that certain regions in the parameter space exist for which only a few fixed points are present. This suddenly opens the door for numerical methods that efficiently obtain all fixed points and enhance the accuracy by combining them.

When defining well-behaved regions of BP we would consequently like to initiate a shift from conditions for uniqueness to conditions for structured, albeit possibly multiple, fixed points. Fairly substantial work has been conducted in this matter for the special case of optimization problems. Developing similar results for more general models is a necessary step to assess the relevance of methods that rely on considering multiple fixed points and to elaborate on their capabilities.

- ### 2. Global Convergence Properties:
- We have put a lot of emphasize on the local stability analysis in Chapter 5. Let us stress the notion of *local* here again: we considered a fixed point as stable if the messages will converge to the fixed point inside its neighborhood. In general, one can, however, not count on bringing the messages close enough to the fixed point and it is not only the local stability that is important for the convergence properties of BP; even more so it is the *region of attraction*. The region of attraction describes the range of message values for which BP will converge to a given fixed point. This is particularly relevant in practice as BP sometimes fail to converge despite the existence of stable fixed points (unless initialized very close to the fixed point).

Such a global perspective also highlights the influence of the initial message values on the convergence properties of BP. Besides being convenient, uniform initialization tends to favor good fixed points; it remains an open problem to answer what is so special about uniform initialization and under which conditions uniform initialization is optimal. We expect that extending the local stability analysis and accounting for the regions of attraction as well will provide many answers to those questions. Besides, this could potentially lead to variants of BP that, even if initialized far away from any fixed point, explore the

parameter space in a way such that the messages enter a region of attraction at some point, ultimately driving the messages into convergence.

3. **Enhancing BP with Respect to Marginal Accuracy:** The main contribution of Chapter 7 was to elaborate on the difference between the accuracy of the marginals and the partition function. Our analysis revealed that the fixed point providing the most accurate marginals may only be a local minimum. Note that most variants that aim to enhance BP are either heuristics that only aim to enforce convergence or are theoretically well-motivated and focus on the quality of the partition function approximation.

For many applications, we are, however, ultimately interested in enhancing the quality of the marginals (cf. error-correcting codes in Section 5.9). It would thus be of great relevance to manipulate the update equation in a way that enforces convergence towards a fixed point with accurate marginals. Enforcing convergence towards one particular fixed point would further enable one to come up with realistic performance guarantees, even in the presence of multiple fixed points.

We have proposed SBP that is guaranteed to obtain the most accurate marginals for attractive models with unidirectional potentials. Looking at the evolution of the fixed points on patch-potential models, it seems as if SBP is capable of tracking the state-preserving fixed point as well. Thus, SBP is already one possible candidate for obtaining accurate marginals. It would be interesting to elaborate on this observation and substantiate this claim in a more formal way. Essentially, the advantages of SBP boil down to providing a favorable initialization for BP. This highlights the importance of the initial messages once more and raises the question if there is a general way of initializing the messages that will lead to accurate marginals.

Besides the role of the initialization, scheduling plays an important role in determining the fixed point of BP. Again, one would hope that “good” scheduling methods enforce convergence towards an accurate fixed point. Existing scheduling methods are either adaptive or only take the structure of the model into account. The performance of BP, however, not only depends on the structure of the model but on the potentials as well. Therefore, if the aim is to obtain accurate marginals, the value of the potentials should not be neglected and it seems promising to come up with a fixed schedule that takes the values of the potentials into account as well, for example by scheduling the messages in a way that prioritizes regions with strong local potentials.

A

Appendix

A.1 Proofs from Chapter 6

This Section contains all the detailed proofs deferred from Section 6.4.

A.1.1 Proof of Lemma 4

Essentially, we first show that $\mu_{ij}(X_j = 1)/\mu_{ij}(X_j = -1)$ increases monotonically with J_{ij} and then express the pseudomarginals in terms of (3.2) and (3.3).

Let us denote the messages on \mathcal{U}_0 and on \mathcal{U}_1 by $\mu_{ij}(x_j)_0$ and $\mu_{ij}(x_j)_1$ respectively. Further, let all local potentials be positive, i.e., $\theta_i > 0$, and let all pairwise potentials of \mathcal{U}_1 be ϵ -larger than those of \mathcal{U}_0 , i.e., $0 < J_{ij}^0 = J_{ij}^1 - \epsilon$. Note that by assumption $m_i \in (0, 1]$ so that

$$\mu_{ij}(X_j = +1) \geq \mu_{ij}(X_j = -1). \quad (\text{A.1})$$

First, we show that for all $(i, j) \in \mathbf{E}$

$$\frac{\mu_{ij}^\circ(X_j = +1)_0}{\mu_{ij}^\circ(X_j = -1)_0} < \frac{\mu_{ij}^\circ(X_j = +1)_1}{\mu_{ij}^\circ(X_j = -1)_1}. \quad (\text{A.2})$$

Therefore, consider the update rule of (3.1) for both states

$$\mu_{ij}^{n+1}(X_j = +1)_1 \propto e^{J_{ij} + \theta_i + \epsilon} \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(X_i = +1)_1 + e^{-J_{ij} - \theta_i - \epsilon} \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(X_i = -1)_1, \quad (\text{A.3})$$

and

$$\mu_{ij}^{n+1}(X_j = -1)_1 \propto e^{-J_{ij} + \theta_i - \epsilon} \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(X_i = +1)_1 + e^{J_{ij} - \theta_i + \epsilon} \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(X_i = -1)_1. \quad (\text{A.4})$$

In (A.3) the larger product is multiplied by e^ϵ and the smaller product is divided by e^ϵ . For (A.4) it is exactly the other way round so that the ratio between the messages increases which proofs (A.2). We shall denote the imposed difference $\delta \in \mathbb{R}_+^*$ on the messages by

$$\mu_{ij}^\circ(X_j = +1)_1 = \mu_{ij}^\circ(X_j = +1)_0 + \delta, \quad (\text{A.5})$$

$$\mu_{ij}^\circ(X_j = -1)_1 = \mu_{ij}^\circ(X_j = -1)_0 - \delta. \quad (\text{A.6})$$

Second, we show that $m_i^0 < m_i^1$ which is an immediate consequence of plugging (A.5) and (A.6) into (3.2).

Finally, it remains to show that $0 < \chi_{ij}^0 \stackrel{(i)}{<} \chi_{ij}^1$. Without loss of generality we assume that all variables have equal degree $d + 1$ and constant coupling strength $J_{ij} = J$. First we show that (i) holds, i.e., $\chi = \chi_{ij}^0$ is positive. Let us express the marginals by (3.3) and denote the messages

by $\mu = \mu_{ij}(X_j = 1)_0$. It follows that $\mu_{ij}(X_j = -1)_0 = (1 - \mu)$ and that

$$\chi = e^{J+2\theta} \mu^{2d} + e^{J-2\theta} (1-\mu)^{2d} - 2e^{-J} \mu^d (1-\mu)^d. \quad (\text{A.7})$$

Let us further represent the messages by $\mu = 1/2 + x$ with $x \in [0, 1/2]$. It follows that

$$\begin{aligned} \chi &\stackrel{(a)}{\geq} \left((1/2 + x)^{2d} + (1/2 - x)^{2d} \right) - 2(1/2 + x)^d (1/2 - x)^d \\ &= \left((1/2 - x)^d - (1/2 + x)^d \right)^2 \\ &\stackrel{(b)}{\geq} 0, \end{aligned} \quad (\text{A.8})$$

where (a) follows from neglecting all exponential terms and thus upper bounding the positive term and lower bounding the negative term (with equality if and only if $J = 0$ and $\theta = 0$) and (b) is a direct consequence of the square in (A.8). Now let us show that (ii) holds, i.e., χ increases monotonically, by taking the derivative of (A.7), so that

$$\begin{aligned} \frac{\partial}{\partial \mu} \chi &= 2d \left(e^{J+2\theta} \mu^{2d-1} - e^{J-2\theta} (1-\mu)^{2d-1} \right) + 2de^{-J} \left(\mu^d (1-\mu)^{d-1} - \mu^{d-1} (1-\mu)^d \right) \quad (\text{A.9}) \\ &\stackrel{(a)}{\geq} 2de^{-J} \left(\mu^d (1-\mu)^{d-1} - \mu^{d-1} (1-\mu)^d \right) \\ &\stackrel{(b)}{\geq} 0, \end{aligned} \quad (\text{A.10})$$

where (a) follows from neglecting the, strictly positive, first term in (A.9), and (b) is a direct consequence from (A.1).

A.1.2 Proof of Theorem 9

A unique start point $\tilde{P}_B^m(\zeta = 0)$ exists by Theorem 6 that equals the exact pseudomarginals $P_B(\zeta = 0)$ and, for $\theta_i > 0$, has positive mean and correlation), i.e., $m_i(\zeta = 0) > 0$.

Consequently, Lemma 4 applies, which further implies that $m_i(\zeta)$ and $\chi_{ij}(\zeta)$ are monotonically increasing; moreover, $m_i(\zeta)$ and $\chi_{ij}(\zeta)$ are continuous by Theorem 7. This further implies that the Bethe free energy $\mathcal{F}_B^m(\zeta)$ decreases. Let $\mathcal{F}_B^m(\zeta = 1)$ correspond to the endpoint of the solution path $c(\zeta)$ that emerges from the origin; then, it immediately follows that the error with respect to the endpoint $\mathcal{F}_B^m(\zeta = 1)$ decreases along the solution path: i.e., consider two arbitrary values $m, k \in [0, 1]$ such that $k > m$, then $|\mathcal{F}_B^m(\zeta_m) - \mathcal{F}_B^m(\zeta = 1)| \geq |\mathcal{F}_B^m(\zeta_k) - \mathcal{F}_B^m(\zeta = 1)|$.

It remains to show that SBP obtains the fixed point $\mathcal{F}_B^m(\zeta = 1)$ that minimizes the error with respect to the exact free energy, i.e., $m = \arg \min_{m \in \mathbf{M}} E_{\mathcal{Z}}(m)$. Therefore consider the fact, that attractive models with $\theta_i > 0$ have a unique fixed point that satisfies $m_i(\zeta) \in (0, 1]$ [YFW05]. A second minimum with negative means, however, may emerge for sufficiently large values of J_{ij} . We denote this alternative stationary point by \mathcal{F}_B^n . This minimum \mathcal{F}_B^n , if it exists, is close to being symmetric, i.e., $m_i^m(\zeta) - \epsilon = -m_i^n(\zeta)$ and $\chi_{ij}^m(\zeta) = \chi_{ij}^n(\zeta) + \epsilon$.

Now let us express the Bethe free energy in (3.24) of both fixed points \tilde{P}_B^m and \tilde{P}_B^n in terms of their energy and entropy according to $\mathcal{F}_B = E_B - S_B$ (cf. Section 3.3). Then, as a consequence of symmetry of the entropy $S_B(\tilde{P}_B^n) \cong S_B(\tilde{P}_B^m)$ and as a consequence of singleton marginals that are not aligned to the local potentials in (3.21) $E_B(\tilde{P}_B^n) > E_B(\tilde{P}_B^m)$. It follows that $\mathcal{F}_B(\tilde{P}_B^n) \geq \mathcal{F}_B(\tilde{P}_B^m)$ (cf. [PAM11]). Consequently, with $\mathcal{F}_B(\tilde{P}_B^m)$ being more negative it follows that the fixed point m constitutes the global minimum of the Bethe free energy.

That is, SBP proceeds along a solution path that leads towards the global minimum of the

Bethe approximation. In particular, by considering the fact that the exact free energy is upper bounded by the Bethe approximation for attractive models [Ruo13], this implies that the fixed point obtained by SBP indeed minimizes the approximation error, i.e., $m = \arg \min_{m \in \mathbf{M}} E_Z(m)$ holds.

This concurrently implies that $\tilde{P}_B^m(\zeta)$ is optimal with respect to marginal accuracy, i.e., no stable fixed point – which corresponds to a local minimum of the Bethe free energy – exists that provides more accurate marginals. Note that attractive models exhibit so-called replica symmetric solutions where the Bethe free energy has two minima at most. In particular, this allows one to express the exact marginals as a convex combination of all fixed points, i.e., $P_B = \frac{1}{\sum_{z_B \in \mathbf{M}} z_B} \sum_{(\tilde{P}_B, z_B) \in \mathbf{M}} \tilde{P}_B \cdot z_B$ (cf. Section 3.5.4).

It follows by the existence of at most two solutions that the fixed point that minimizes the Bethe free energy is also more accurate.

A.2 Proofs from Chapter 7

This Section contains all the detailed proofs for Chapter 7

A.2.1 Proof of Theorem 11

Here we show how all incoming messages from outside the patch can be subsumed into an effective local field.

Proof. First let us revisit the update equation from X_i to $X_j : X_j \in \mathbf{X}_i$

$$\mu_{ij}^{n+1}(x_j) \propto \sum_{x_i \in \mathcal{X}} \Phi(x_i, x_j) \Phi(x_i) \prod_{X_k \in \{\partial(i) \setminus X_j\}} \mu_{ki}^n(x_i).$$

Now we group the incoming messages into two groups, i.e., messages coming from outside the patch and messages coming from inside the patch so that

$$\mu_{ij}^{n+1}(x_j) \propto \sum_{x_i \in \mathcal{X}} \Phi(x_i, x_j) \Phi(x_i) \prod_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} \mu_{ki}^n(x_i) \prod_{X_k \in \{\mathbf{X}_i \cap \partial(i) \setminus X_j\}} \mu_{ki}^n(x_i).$$

Now we make use of the fact that we are only dealing with binary random variables for which $\mu_{ij}(X_j = -1) = (1 - \mu_{ij}(X_j = 1))$ and express the message explicitly by

$$\begin{aligned} \mu_{ij}^{n+1}(x_j) &\propto \exp(Jx_j) \exp(\theta_i) \prod_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} \mu_{ki}^n(X_i = 1) \prod_{X_k \in \{\mathbf{X}_i \cap \partial(i) \setminus X_j\}} \mu_{ki}^n(X_i = 1) \\ &\quad + \exp(-Jx_j) \exp(-\theta_i) \prod_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} (1 - \mu_{ki}^n(X_i = 1)) \prod_{X_k \in \{\mathbf{X}_i \cap \partial(i) \setminus X_j\}} (1 - \mu_{ki}^n(X_i = 1)). \end{aligned} \quad (\text{A.11})$$

We now want to get rid of the first product and absorb the influence of these messages into the local field. In particular we aim to express it according to

$$\begin{aligned} \exp(\tilde{\theta}_i x_i) &= \Phi(x_i) \prod_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} \mu_{ki}^n(x_i) \\ &= \exp((\theta_i + c)x_i) \cdot \exp(g). \end{aligned} \quad (\text{A.12})$$

In order to do so we take the logarithm of the product over all messages in (A.12) and put them into the exponent with the local field so that for $x_i = +1$

$$\begin{aligned} \exp(\tilde{\theta}_i) &= \exp\left(\theta_i + \sum_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} \log(\mu_{ki}^n(X_i = 1))\right) \\ &= \exp\left(\theta_i + \sum_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} c_i + \sum_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} g_i\right), \end{aligned} \quad (\text{A.13})$$

and for $x_i = -1$

$$\begin{aligned} \exp(-\tilde{\theta}_i) &= \exp\left(-\theta_i + \sum_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} \log(1 - (\mu_{ki}^n(X_i = 1)))\right) \\ &= \exp\left(-\theta_i - \sum_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} c_i + \sum_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} g_i\right). \end{aligned} \quad (\text{A.14})$$

Equating the coefficients for c_i and g_i in (A.13) and (A.14) gives us the final results:

$$c_i = \frac{1}{2} \left(\log \mu_{ki}(X_i = 1) - \log(1 - \mu_{ki}(X_i = 1)) \right),$$

$$c = \sum_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} \operatorname{atanh} \left(2\mu_{ki}(X_i = 1) - 1 \right), \quad (\text{A.15})$$

$$g = \frac{1}{2} \log \prod_{X_k \in \partial(i) \setminus \{X_j \cap \mathbf{X}_i\}} \left(\mu_{ki}(X_i = 1) - \mu_{ki}(X_i = 1)^2 \right). \quad (\text{A.16})$$

Note that we can further express the message for the second state $\mu_{ij}(X_j = -1)$ in a similar way, with the only difference that the values of the pairwise potentials change. Consequently we get exactly the same result for g again; this allows us to neglect the influence of g altogether, as it will be canceled out when normalizing the messages so that the sum up to one. \square

A.2.2 Proof of Corollary 15.1

We want to compare error of the state-preserving fixed point $E_P(p)$ to the error $E_P(q)$ of a fixed point that has all marginals biased towards one state. We already discussed the error-ratio in a general manner in the proof of Theorem 15 in Section 7.4.3; here we provide a more accessible proof for the special case of a model with two patches, i.e., for Example 10.

Proof. For a symmetric model with two equal-sized patches we evaluate the error ratio between the state-preserving fixed point p and one of the fixed points that have all marginals biased towards one state, these are q and r and have *symmetric* marginals. Further assume that the fixed points q and r minimize the Bethe free energy, i.e.,

$$\mathcal{F}_B^q = \mathcal{F}_B^r < \mathcal{F}_B^p, \quad (\text{A.17})$$

$$\mathcal{Z}_B^q = \mathcal{Z}_B^r > \mathcal{Z}_B^p. \quad (\text{A.18})$$

Then we want to show that (A.18) does not imply that $E_P(q) < E_P(p)$, i.e., we want to show that $\frac{E_P(p)}{E_P(q)} < 1$ despite (A.18); therefore, we express the ratio of the marginal errors according to

$$\begin{aligned} \frac{E_P(p)}{E_P(q)} &= \frac{\sum_{X_i} |\sum_{m \setminus p} \mathcal{Z}_B^m Q_i(m, p)|^2}{\sum_{X_i} |\sum_{m \setminus q} \mathcal{Z}_B^m Q_i(m, q)|^2} \\ &= \frac{\sum_{X_i} |\mathcal{Z}_B^q Q_i(q, p) + \mathcal{Z}_B^r Q_i(r, p)|^2}{\sum_{X_i} |\mathcal{Z}_B^q Q_i(r, q) + \mathcal{Z}_B^p Q_i(p, q)|^2}. \end{aligned} \quad (\text{A.19})$$

Note that because of (A.18) we have

$$\begin{aligned} \frac{E_P(p)}{E_P(q)} &= \frac{\sum_{X_i} |\mathcal{Z}_B^q (Q_i(q, p) + Q_i(r, p))|^2}{\sum_{X_i} |\mathcal{Z}_B^q Q_i(r, q) + \mathcal{Z}_B^p Q_i(p, q)|^2} \\ &\stackrel{(a)}{=} \frac{\sum_{X_i} |\mathcal{Z}_B^q (Q_i(q, p) + Q_i(r, p))|^2}{\sum_{X_i} (\mathcal{Z}_B^q Q_i(q, r) + \mathcal{Z}_B^p Q_i(q, p))^2}. \end{aligned} \quad (\text{A.20})$$

where (a) follows from the fact that $Q_i(r, q) < 0$, $Q_i(p, q) < 0$, and the symmetry property (7.17). We further denote the constant difference between the biased fixed points by

$$0 < Q_i(q, r) = d < 1. \quad (\text{A.21})$$

We can use the expansion property (7.18) to bound the numerator as $|Q_i(q, p) + Q_i(r, p)|^2 < Q_i(q, r)^2 = d^2$ so that

$$\frac{E_P(p)}{E_P(q)} < \frac{\sum_{X_i} Z_B^q d^2}{\sum_{X_i} (Z_B^q d + Z_B^p Q_i(q, p))^2}. \quad (\text{A.22})$$

Which completes the proof as $Q_i(q, p) > 0$ □

A.2.3 Proof of Theorem 16

Let us consider three different stationary points \mathcal{F}_B^p (state-preserving), \mathcal{F}_B^q (biased to one state), and \mathcal{F}_B^m that has some patches flipped. Note that we consider \mathcal{F}_B^q as a limiting case for \mathcal{F}_B^m . We will denote the number of variables that are aligned with the local field N_c and the number of flipped variables N_f .

The set of boundary edges that connects two patches is denoted by

$$\mathbf{E}_P = \{(i, j) \in \mathbf{E} : X_i \in \mathbf{X}_i, X_j \in \mathbf{X}_j \neq \mathbf{X}_i\}, \quad (\text{A.23})$$

and the set of edges that connects two patches that have their variables not aligned is denoted by

$$\mathbf{E}_C = \{(i, j) \in \mathbf{E} : X_i \in \mathbf{X}_i, X_j \in \mathbf{X}_j \neq \mathbf{X}_i, \text{sgn}(\tilde{P}_{X_i}(X_i = 1) - 0.5) \neq \text{sgn}(\tilde{P}_{X_j}(X_j = 1) - 0.5)\}. \quad (\text{A.24})$$

Note that \mathbf{E}_P is constant for a specified model, whereas \mathbf{E}_C depends on the specific fixed point. We consequently have $\mathbf{E}_C \leq \mathbf{E}_P$ with equality for the state preserving fixed point p .

Our analysis is restricted to $(\theta, J) \in (II)$ per definition: one crucial consequence is that $J > J_A(\theta)$ and that most marginals either have $\tilde{P}_{X_i}(x_i) \approx 1$ or $\tilde{P}_{X_i}(x_i) \approx 0$. We will exploit this fact and express all marginals according to $\tilde{P}_{X_i}(x_i) \in \{0, 1\}$ which allows us to simplify \mathcal{F}_B , as defined in Section 3.3, according to

$$\mathcal{F}_B^p = -N\theta - (|\mathbf{E}| - 2|\mathbf{E}_P|)J - S_B^p \quad (\text{A.25})$$

$$\mathcal{F}_B^m = -(N_c - N_f)\theta - (|\mathbf{E}| - 2|\mathbf{E}_C|)J - S_B^m. \quad (\text{A.26})$$

Let $\Delta S_B = S_B^p - S_B^m$ be the difference in the entropy, then we can express the conditions for the state-preserving fixed point to have a lower value $\mathcal{F}_B^p \leq \mathcal{F}_B^m$ according to

$$\begin{aligned} -N\theta - (|\mathbf{E}| - 2|\mathbf{E}_P|)J &\leq -(N_c - N_f)\theta - (|\mathbf{E}| - 2|\mathbf{E}_C|)J + \Delta S_B \\ 2J(-|\mathbf{E}_C| + |\mathbf{E}_P|) &\leq \theta(N - N_c + N_f) + \Delta S_B \end{aligned} \quad (\text{A.27})$$

Now let us express (A.27) for the fixed points that has all variables biased to one state, i.e., $|\mathbf{E}_C| = 0$. Then, the state-preserving fixed point has a lower value $\mathcal{F}_B^p < \mathcal{F}_B^q$ if

$$2J|\mathbf{E}_P| \leq \theta(N - N_c + N_f) + \Delta S_B \quad (\text{A.28})$$

A.3 Pseudocode

We present the pseudocode for NIBP and WDBP. Removing the if then else clause in line 8 to 11 of NIBP and substituting it with $\mu_m^n \leftarrow \mu_m^{n+1}$ reduces Algorithm 3 to RBP. The maximum number of iterations is denoted by $N_{BP} = 2.5 \cdot 10^5$ and $\epsilon = 10^{-3}$. $\text{NrOfMessages} = 2|\mathbf{E}|$ denotes the overall number of messages in the graph.

Algorithm 3: Noise Injection Belief Propagation (NIBP)

input : Graph $G = (\mathbf{X}, \mathbf{E})$
output: Converged messages μ^n

- 1 initialization
- 2 **for** $m \leftarrow 1$ **to** NrOfMessages **do**
- 3 $\mu_m^{n+1} \leftarrow \text{ComputeUpdate}(\mu^n)$
- 4 $r_m \leftarrow |\mu_m^n - \mu_m^{n+1}|$
- 5 $k \leftarrow 1$
- 6 **while** $n < N_{BP}$ **and** $\max |\mu^n - \mu^{n+1}| > \epsilon$ **do**
- 7 $m \leftarrow \arg \max_m r^n$
- 8 **if** $\text{OscillationDetection}(\mu_m^n, L)$ **then**
- 9 $\mu_m^n \leftarrow \mu_m^{n+1} + \mathcal{N}(0, \sigma)$
- 10 **else**
- 11 $\mu_m^n \leftarrow \mu_m^{n+1}$
- 12 **for** $j \leftarrow 1$ **to** NrOfMessages **do**
- 13 $\mu_j^{n+1} \leftarrow \text{ComputeUpdate}(\mu^n)$
- 14 $r_j \leftarrow |\mu_j^{n+1} - \mu_j^n|$
- 15 $n = n + 1$

Algorithm 4: Weight Decay Belief Propagation (WDBP)

input : Graph $\mathcal{G} = (\mathbf{X}, \mathbf{E})$
output: Converged messages μ^n

- 1 initialization
- 2 **for** $m \leftarrow 1$ **to** NrOfMessages **do**
- 3 $\mu_m^{n+1} \leftarrow \text{ComputeUpdate}(\mu^n)$
- 4 $r_m \leftarrow |\mu_m^n - \mu_m^{n+1}|$
- 5 $\text{NrUpdates}(m) \leftarrow 1$
- 6 $k \leftarrow 1$
- 7 **while** $n < N_{BP}$ **and** $\max |\mu^n - \mu^{n+1}| > \epsilon$ **do**
- 8 $m \leftarrow \arg \max_m r^n$
- 9 $\mu_m^n \leftarrow \mu_m^{n+1}$
- 10 $\text{NrUpdates}(m) \leftarrow \text{NrUpdates}(m) + 1$
- 11 **for** $j \leftarrow 1$ **to** NrOfMessages **do**
- 12 $\mu_j^{n+1} \leftarrow \text{ComputeUpdate}(\mu^n)$
- 13 $r_j \leftarrow \frac{|\mu_j^{n+1} - \mu_j^n|}{\text{NrUpdates}(j)}$
- 14 $n = n + 1$

B

Guide to Quotes

The attentive reader has probably recognized that this thesis features an opening quote for every chapter. It was important to me to connect with all the quotes on a personal level (be it by one's art or by a general appreciation of a person's life) Yet, this was not enough: every quote must also have a close connection to the respective chapter. This connection, however, may not always be as obvious to the reader as it was to me. It is precisely for this reason that I will now reveal the perceived meaning of each quote, and how I see it connect to the content of the chapter.

Chapter 1: We quote the commonly used English translation "*Uncertainty is an uncomfortable position. But certainty is an absurd one.*" of the original line "*Le doute n'est pas une état bien agréable, mais l'assurance est un état ridicule*" [Vol17, p.703]. This statement from Voltaire elegantly connects to the central concept of this thesis, probabilistic graphical model, as it can be interpreted as a praise for probability.

Chapter 2: "*A mind is like a parachute. It doesn't work if it is not open.*" This quote by Frank Zappa seems to align nicely with a background-chapter that brings different scientific fields together into one powerful concept. In particular, since it was Frank Zappa's ability to draw inspiration from an unimaginable wide range of musical genres that finally coalesced together and created some stunning albums.

Chapter 3: Ludwig Boltzmann played a formative role in the foundational years of statistical physics and laid out the groundwork for many concepts discussed in this chapter. Beyond that, there is undeniable truth in his statement "*The stars bend like slaves to laws not decreed for them by human intelligence, but gleaned from them.*" quoted from [GRE91]. One could argue that this whole thesis adheres to this statement as we put considerable effort into carefully studying the behavior of belief propagation to develop a clearer picture of its underlying working-principles.

Chapter 4: "*Inside a broken clock; Splashing the wine; With all the rain dogs; Taxi, we'd rather walk...*" [Wai85]. The song figuratively refers to the rain dogs as the straying dogs that are seemingly lost after the rain has washed away all their scent. The underlying meaning is elegantly encompassed by just another quote of Tom Waits that states: "*We are buried beneath the weight of information, which is being confused with knowledge; quantity is being confused with abundance and wealth with happiness.*". Although arguable more general and touching a couple of interesting points of today's society, Tom Waits' statement reflects the experiences one makes when acquainting oneself with the subtleties of a new field. When embarking into uncharted territory and trying to grasp all impressions one can only sympathize with the rain dogs that must feel similar.

Chapter 5: "*Travel makes one modest. You see what a tiny place you occupy in the world.*" [Fla96]. Flaubert urges the need for always taking a look at the surrounding to put things into perspective. BP often behaves like a creature of habit, once it decides for one specific fixed point it neglects all others. In this chapter we break up with this habit and take a global perspective on the solution space, essentially traveling through the solution space.

Chapter 6: *“The most significant dimension of freedom is the freedom from one’s own ego - in other words, from the feeling that I am the center of everything.”* Voytek Kurtyka is an outstanding mountaineer that played an important role in bringing the modern alpine-style climbing into the greater ranges. More importantly, he was always honest about his inner struggle between the desire to nourish one’s ego and the awareness that there is no inbred reason to take oneself to serious. One could argue that our proposed algorithm, self-guided BP, undergoes a similar struggle. At the beginning, all variables take up a purely ego-centered standpoint and no interactions between them take place. Only incrementally, as the couplings are incorporate, the variables contribute to the joint assignment of the model and put their own constraints behind.

Chapter 7: *“To the wise, life is a problem; to the fool, a solution.”* Wise and fool are maybe strong words; but one could have been satisfied with the preceding chapters and could have neglected the unanswered questions lingering around instead of being curious and trying to gain some further insights. On a side-note, Marcus Aurelius was a great proponent of the philosophy of stoicism, a philosophy that is arguable beneficial when working on a thesis for many years.

Chapter 8: *“There must be some way out of here; Said the joker to the thief“* [Dyl67]. This song puts you in the midst of a discussion, with the joker struggling with his purpose. Even though it was inevitably to experience some struggles in the process of writing everything up, there is no easy way out and, ultimately, it is just a question of getting things done before one can appreciate the completed task.

Bibliography

- [ACSW18] Sungsoo Ahn, Michael Chertkov, Jinwoo Shin, and Adrian Weller. Gauged mini-bucket elimination for approximate inference. *Proceedings of AISTATS*, 2018.
- [AG03] E. Allgower and K. Georg. *Introduction to Numerical Continuation Methods*. Society for Industrial and Applied Mathematics, 2003.
- [AM00] Srinivas M Aji and Robert J McEliece. The generalized distributive law. *IEEE Transactions on Information Theory*, 46(2):325–343, 2000.
- [AS88] Winfried Auzinger and Hans J Stetter. An elimination algorithm for the computation of all zeros of a system of multivariate polynomial equations. In *Proceedings of International Conference on Numerical Mathematics*, pages 11–30. Springer, 1988.
- [AS89] Winfried Auzinger and Hans J Stetter. A study of numerical elimination for the solution of multivariate polynomial systems. Technical report, Institut für Angewandte und Numerische Mathematik, Vienna University of Technology, 1989.
- [Ber75] David N. Bernstein. The number of roots of a system of equations. *Functional Analysis and Its Applications*, 9:1–4, 1975.
- [Bet35] Hans A Bethe. Statistical theory of superlattices. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences*, 150(871):552–575, 1935.
- [BF11] Richard L Burden and Douglas J Faires. *Numerical Analysis*. Brooks Cole Cengage, 9 edition, 2011.
- [BGT93] Claude Berrou, Alain Glavieux, and Punya Thitimajshima. Near Shannon limit error-correcting coding and decoding: Turbo-codes. 1. In *Proceedings of IEEE International Conference on Communications*, volume 2, pages 1064–1070, 1993.
- [BH11] Andries E Brouwer and Willem H Haemers. *Spectra of Graphs*. Springer, 2011.
- [BHSW] Daniel J. Bates, Jonathan D. Hauenstein, Andrew J. Sommese, and Charles W. Wampler. Bertini: Software for numerical algebraic geometry. Available at bertini.nd.edu with permanent doi: dx.doi.org/10.7274/R0H41PB5.
- [Bin87] K Binder. Finite size effects on phase transitions. *Ferroelectrics*, 73(1):43–67, 1987.
- [Bis06] Christopher Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [BKMZ07] Alfredo Braunstein, Farbod Kayhan, Guido Montorsi, and Riccardo Zecchina. Encoding for the blackwell channel with reinforced belief propagation. In *IEEE International Symposium on Information Theory*, 2007.
- [BKvdEvdG01] Hans L Bodlaender, Arie MCA Koster, Frank van den Eijkhof, and Linda C van der Gaag. Pre-processing for triangulation of probabilistic networks. In *Proceedings of UAI*, pages 32–39, 2001.

- [BMZ05] A Braunstein, M Mézard, and R Zecchina. Survey propagation: An algorithm for satisfiability. *Random Structures & Algorithms*, 27(2):201–226, 2005.
- [Bru67] Stephen G Brush. History of the Lenz-Ising model. *Reviews of Modern Physics*, 39(4):883, 1967.
- [Buc65] Bruno Buchberger. *Ein Algorithmus zum Auffinden der Basiselemente des Restklassenringes nach einem nulldimensionalen Polynomideal*. PhD thesis, University of Innsbruck, 1965.
- [Buc06] Bruno Buchberger. Bruno buchberger’s phd thesis 1965: An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. *Journal of Symbolic Computation*, 41(3-4):475–511, 2006.
- [BYGRS12] D Batra, P Yadollahpour, A Guzman-Rivera, and G Shakhnarovich. Diverse m-best solutions in markov random fields. In *Computer Vision–ECCV 2012*, pages 1–16. Springer, 2012.
- [CC06] Michael Chertkov and Vladimir Y Chernyak. Loop series for discrete statistical models on graphs. *Journal of Statistical Mechanics: Theory and Experiment*, (06), 2006.
- [CCG⁺11] Venkat Chandrasekaran, Misha Chertkov, David Gamarnik, Devavrat Shah, and Jinwoo Shin. Counting independent sets using the Bethe approximation. *SIAM Journal on Discrete Mathematics*, 25(2):1012–1034, 2011.
- [CCT08] Michael Chertkov, Vladimir Y Chernyak, and Razvan Teodorescu. Belief propagation and loop series on planar graphs. *Journal of Statistical Mechanics: Theory and Experiment*, (05), 2008.
- [CL15] Tianran Chen and Tien-Yien Li. Homotopy continuation method for solving systems of nonlinear and polynomial equations. *Communications in Information and System*, 15(2):119–307, 2015.
- [CLL14] Tianran Chen, Tsung-Lin Lee, and Tien-Yien Li. Hom4ps-3: A parallel numerical solver for systems of polynomial equations based on polyhedral homotopy continuation methods. In *Proceedings of International Congress on Mathematical Software*, pages 183–190. Springer, 2014.
- [CLO92] D Cox, J Little, and D O’Shea. *Ideals, Varieties, and Algorithms*, volume 3. Springer, 1992.
- [CLO05] D Cox, J Little, and D O’Shea. *Using Algebraic Geometry*, volume 185. Springer, 2005.
- [CM17] T Chen and D Mehta. On the network topology dependent solution count of the algebraic load flow equations. *IEEE Transactions on Power Systems*, 2017.
- [CMN16] T Chen, D Mehta, and M Niemerg. A network topology dependent upper bound on the number of equilibria of the Kuramoto model. *arXiv preprint arXiv:1603.05905*, 2016.
- [Coo90] Gregory F Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2-3):393–405, 1990.
- [COP19] Amin Coja-Oghlan and Will Perkins. Bethe states of random factor graphs. *Communications in Mathematical Physics*, 366(1):173–201, 2019.

- [DE05] Alicia Dickenstein and Ioannis Z Emiris. *Solving Polynomial Equations*. Springer, 2005.
- [Die16] Reinhard Diestel. *Graph Theory*. Springer, 2016.
- [DL93] Paul Dagum and Michael Luby. Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artificial Intelligence*, 60(1):141–153, 1993.
- [DM10] Amir Dembo and Andrea Montanari. Ising models on locally tree-like graphs. *The Annals of Applied Probability*, 20(2):565–592, 2010.
- [DR03] Rina Dechter and Irina Rish. Mini-buckets: A general scheme for bounded inference. *Journal of the ACM*, 50(2):107–153, 2003.
- [Dyl67] Bob Dylan. *All Along The Watchtower*. Columbia Records, 1967.
- [EG13] Frederik Eaton and Zoubin Ghahramani. Model reductions for inference: Generality of pairwise, binary, and planar factor graphs. *Neural Computation*, 25(5):1213–1260, 2013.
- [EH16] Bradley Efron and Trevor Hastie. *Computer Age Statistical Inference*. Cambridge University Press, 2016.
- [EMK06] Gal Elidan, Ian McGraw, and Daphne Koller. Residual belief propagation: Informed scheduling for asynchronous message passing. In *Proceedings of UAI*, 2006.
- [Erd11] László Erdős. Universality of wigner random matrices: a survey of recent results. *Russian Mathematical Surveys*, 66, 2011.
- [Fey72] Richard P. Feynman. *Statistical Mechanics: A Set Of Lectures*. Benjamin/Cummings, 1972.
- [Fis67] Michael E Fisher. Critical temperatures of anisotropic ising lattices. ii. general upper bounds. *Physical Review*, 162(2):480, 1967.
- [Fla96] Gustave Flaubert. *Flaubert in Egypt: A Sensibility on Tour: A Narrative Drawn from Gustave Flaubert's Travel Notes & Letters*. Penguin, 1996.
- [Gal68] Robert G Gallager. *Information Theory and Reliable Communication*, volume 2. Springer, 1968.
- [Geo11] Hans O Georgii. *Gibbs Measures and Phase Transitions*, volume 9. Walter de Gruyter, 2011.
- [GJ07] Amir Globerson and Tommi S Jaakkola. Convergent propagation algorithms via oriented trees. In *Proceedings of UAI*, 2007.
- [GMK07] Vicenç Gómez, Joris M Mooij, and Hilbert J Kappen. Truncating the loop series expansion for belief propagation. *Journal of Machine Learning Research*, 8, 2007.
- [GRE91] GEORGE GREENSTEIN. Science: The bulldog: A profile of ludwig boltzmann. *The American Scholar*, 60(1):97–105, 1991.
- [Gri67] Robert B Griffiths. Correlations in Ising ferromagnets. ii. external magnetic fields. *Journal of Mathematical Physics*, 8(3):484–489, 1967.

- [Gro59] David M Grobman. Homeomorphism of systems of differential equations. *Doklady Akademii Nauk SSSR*, 128(5):880–881, 1959.
- [Gé77] Toulouse Gérard. Theory of the frustration effect in spin glasses. *Communications on Physics*, 2:115–119, 1977.
- [Har60] Philip Hartman. A lemma in the theory of structural stability of differential equations. *Proceedings of the American Mathematical Society*, 11(4):610–620, 1960.
- [HC71] John M Hammersley and Peter Clifford. Markov fields on finite graphs and lattices. *Unpublished manuscript*, 1971.
- [Hes03] Tom Heskes. Stable fixed points of loopy belief propagation are minima of the Bethe free energy. In *Proceedings of NIPS*, volume 15, pages 359–366, 2003.
- [Hes04] Tom Heskes. On the uniqueness of loopy belief propagation fixed points. *Neural Computation*, 16(11), 2004.
- [Hop82] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
- [HS86] Geoffrey E Hinton and Terrence J Sejnowski. Learning and relearning in boltzmann machines. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, 1(282-317):2, 1986.
- [HS95] Birkett Huber and Bernd Sturmfels. A polyhedral method for solving sparse polynomial systems. *Mathematics of Computation*, 64(212):1541–1555, 1995.
- [HS08] Tamir Hazan and Amnon Shashua. Convergent message-passing algorithms for inference over general graphs with convex free energies. In *Proceedings of UAI*, 2008.
- [Hua63] Kerson Huang. *Statistical Mechanics*. John Wiley & Sons Inc, 1963.
- [IFW05] Alexander T Ihler, John W Fisher, and Alan S Willsky. Loopy belief propagation: Convergence and effects of message errors. In *Journal of Machine Learning Research*, pages 905–936, 2005.
- [Ihl] Alexander T Ihler. personal communication, UAI, 2017-08-12.
- [Ihl05] Alexander T Ihler. *Inference in Sensor Networks: Graphical Models and Particle Methods*. PhD thesis, MIT, 2005.
- [Ihl07] Alexander T Ihler. Accuracy bounds for belief propagation. In *Proceedings of UAI*, pages 183–190, July 2007.
- [JGJS99] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.
- [JJ97] Tommi Jaakkola and Michael I Jordan. Recursive algorithms for approximating probabilities in graphical models. In *Proceedings of NIPS*, pages 487–493, 1997.
- [Kal60] Rudolph E Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

- [KF98] Frank R. Kschischang and Brendan J. Frey. Iterative decoding of compound codes by probability propagation in graphical models. *IEEE Journal on Selected Areas in Communications*, 16(2):219–230, 1998.
- [KF09] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.
- [KFD⁺07] Daphne Koller, Nir Friedman, Sašo Džeroski, Charles Sutton, Andrew McCallum, Avi Pfeffer, Pieter Abbeel, Ming-Fai Wong, David Heckerman, Chris Meek, et al. *Introduction to Statistical Relational Learning*. 2007.
- [KFL01] Frank R Kschischang, Branden J Frey, and Hans A Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, 47(2):498–519, 2001.
- [Kho78] Askold G Khovanski. Newton polyhedra and the genus of complete intersections. *Funkts. Anal. Pril.*, 12(1):51–61, 1978.
- [KKP18] Christian Knoll, Florian Kulmer, and Franz Pernkopf. Self-guided belief propagation—a homotopy continuation method. *arXiv preprint arXiv:1812.01339*, 2018.
- [KKR05] Achim Kehrein, Martin Kreuzer, and Lorenzo Robbiano. An algebraist’s view on border bases. In *Solving Polynomial Equations*, pages 169–202. Springer, 2005.
- [KMCP18] Christian Knoll, Dhagash Mehta, Tianran Chen, and Franz Pernkopf. Fixed points of belief propagation – an analysis via polynomial homotopy continuation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [Kol06] Vladimir Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 28(10):1568–1583, 2006.
- [KP17] Christian Knoll and Franz Pernkopf. On loopy belief propagation – local stability analysis for non-vanishing fields. In *Proceedings of UAI*, 2017.
- [KP19] Christian Knoll and Franz Pernkopf. Belief propagation: Accurate marginals or accurate partition function – where is the difference? In *Proceedings of UAI*, 2019.
- [KPMC16] Christian Knoll, Franz Pernkopf, Dhagash Mehta, and Tianran Chen. Fixed point solutions of belief propagation. In *NIPS-Workshop: Advances in Approximate Bayesian Inference*, 2016.
- [KR00] Martin Kreuzer and Lorenzo Robbiano. *Computational Commutative Algebra 1*, volume 1. Springer, 2000.
- [KR05] Martin Kreuzer and Lorenzo Robbiano. *Computational Commutative Algebra 2*, volume 2. Springer, 2005.
- [KRTP15] Christian Knoll, Michael Rath, Sebastian Tschatschek, and Franz Pernkopf. Message scheduling methods for belief propagation. In *Proceedings of ECML PKDD*, pages 295–310. Springer, 2015.
- [Kus76] AG Kushnirenko. Newton polytopes and the Bezout theorem. *Funkts. Anal. Pril.*, 10(3), 1976.

- [KV05] Bernhard Korte and Jens Vygen. *Combinatorial Optimization*, volume 3. Springer, 2005.
- [LCMRTR13] Alejandro Lage-Castellanos, Roberto Mulet, Federico Ricci-Tersenghi, and Tommaso Rizzo. Replica cluster variational method: the replica symmetric solution for the 2d random bond ising model. *Journal of Physics A: Mathematical and Theoretical*, 46, 2013.
- [Li97] Tien-Yien Li. Numerical solution of multivariate polynomial systems by homotopy continuation methods. *Acta numerica*, 6:399–436, 1997.
- [Li03] Tien-Yien Li. Solving polynomial systems by the homotopy continuation method. *Handbook of Numerical Analysis*, 11:209–304, 2003.
- [LI11] Qiang Liu and Alexander T Ihler. Bounding the partition function using Hölder’s inequality. In *Proceedings of ICML*, pages 849–856, 2011.
- [LK03] Martijn Leisink and Hilbert J Kappen. Bound propagation. *Journal of Artificial Intelligence Research*, 19:139–154, 2003.
- [Loe04] Hans A Loeliger. An introduction to factor graphs. *IEEE Signal Processing Magazine*, 21(1):28–41, 2004.
- [LS88] Steffen L Lauritzen and David J Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society. Series B*, pages 157–224, 1988.
- [Mac01] David JC MacKay. A conversation about the Bethe free energy and sum-product. *Technical Report of MERL.*, 2001.
- [Mac03] David JC MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [MGW09] Talya Meltzer, Amir Globerson, and Yair Weiss. Convergent message passing algorithms: a unifying view. In *Proceedings of UAI*, pages 393–401, 2009.
- [MJGF09] Ofer Meshi, Ariel Jaimovich, Amir Globerson, and Nir Friedman. Convexifying the Bethe free energy. In *Proceedings of UAI*. AUAI Press, 2009.
- [MK05] Joris M Mooij and Hilbert J Kappen. On the properties of the Bethe approximation and loopy belief propagation on binary networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(11):P11012, 2005.
- [MK07] Joris M Mooij and Hilbert J Kappen. Sufficient conditions for convergence of the sum-product algorithm. *IEEE Transactions on Information Theory*, 53(12):4422–4437, 2007.
- [MK09] Joris M Mooij and Hilbert J Kappen. Bounds on marginal probability distributions. In *Proceedings of NIPS*, pages 1105–1112, 2009.
- [MLF11] Victorin Martin, Jean-Marc Lasgouttes, and Cyril Furtlehner. The role of normalization in the belief propagation algorithm. *arXiv preprint arXiv:1101.4170*, 2011.
- [MM82] Ernst W Mayr and Albert R Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in Mathematics*, 46(3):305–329, 1982.

- [MM84] H Michael Möller and Ferdinando Mora. Upper and lower bounds for the degree of Gröbner bases. *EUROSAM 84*, pages 172–183, 1984.
- [MM09] Marc Mezard and Andrea Montanari. *Information, Physics, and Computation*. Oxford University Press, 2009.
- [Möl93] H Michael Möller. Systems of algebraic equations solved by means of endomorphisms. In *International Symposium on Applied Algebra, Algebraic Algorithms, and Error-Correcting Codes*, pages 43–56. Springer, 1993.
- [Moo08] Joris Marten Mooij. *Understanding and Improving Belief Propagation*. PhD thesis, Radboud University Nijmegen, 2008.
- [Mou07] Bernard Mourrain. Pythagore’s dilemma, symbolic-numeric computation, and the border basis method. In *Symbolic-Numeric Computation*, pages 223–243. Springer, 2007.
- [MPV87] Marc Mezard, Giorgio Parisi, and Miguel Virasoro. *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Co Inc, 1987.
- [MS95] H Michael Möller and Hans J Stetter. Multivariate polynomial equations with multiple zeros solved by matrix eigenproblems. *Numerische Mathematik*, 70(3):311–329, 1995.
- [Mur12] Kevin P Murphy. *Machine Learning: A Probabilistic Perspective*. MIT press, 2012.
- [MWJ99] Kevin P Murphy, Yair Weiss, and Michael I Jordan. Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of UAI*, pages 467–475, 1999.
- [OW01] Manfred Opper and Ole Winther. Tractable approximations for probabilistic models: the adaptive Thouless-Anderson-Palmer mean field approach. *Physical Review Letters*, 86(17):3695, 2001.
- [PA02] Payam Pakzad and Venkat Anantharam. Belief propagation and statistical physics. In *Proceedings of Conference on Information Sciences and Systems*, 2002.
- [PAM11] Xaq Pitkow, Yashar Ahmadian, and Ken D Miller. Learning unbelievable probabilities. In *Proceedings of NIPS*, pages 738–746, 2011.
- [Pea88] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, 1988.
- [Pei36] R Peierls. Statistical theory of superlattices with unequal concentrations of the components. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences*, 154(881):207–222, 1936.
- [PI17] Wei Ping and Alexander Ihler. Belief propagation in conditional rbms for structured prediction. *Proceedings of AISTATS*, 2017.
- [PPT14] F Pernkopf, R Peharz, and S Tschitschek. *Introduction to Probabilistic Graphical Models*. Academic Press’ Library in Signal Processing, 2014.
- [RG14] Siamak Ravanbakhsh and Russell Greiner. Revisiting algebra and complexity of inference in graphical models. *arXiv preprint arXiv:1409.7410*, 2014.

- [RG15] Siamak Ravanbakhsh and Russell Greiner. Perturbed message passing for constraint satisfaction problems. *Journal of Machine Learning Research*, 16:1249–1274, 2015.
- [RKDW10] Björn S Ruffer, Christopher M Kellett, Peter M Dower, and Steven R Weller. Belief propagation as a dynamical system: The linear case and open problems. *IET Control Theory & Applications*, 4(7):1188–1200, 2010.
- [Rot96] Dan Roth. On the hardness of approximate reasoning. *Artificial Intelligence*, 82(1-2):273–302, 1996.
- [Ruo12] Nicholas Ruozi. The Bethe partition function of log-supermodular graphical models. In *Proceedings of NIPS*, pages 117–125, 2012.
- [Ruo13] Nicholas Ruozi. Beyond log-supermodularity: lower bounds and the Bethe partition function. In *Proceedings of UAI*, pages 546–555, 2013.
- [Sch00] Edward R. Scheinerman. *Invitation to Dynamical Systems*. Prentice Hall, 2000.
- [Shi12] Jinwoo Shin. Complexity of Bethe approximation. In *Proceedings of AISTATS*, pages 1037–1045, 2012.
- [SJ08] David Sontag and Tommi S Jaakkola. New outer bounds on the marginal polytope. In *Proceedings of NIPS*, pages 1393–1400, 2008.
- [SKZ14] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Spectral clustering of graphs with the Bethe Hessian. In *Proceedings of NIPS*, pages 406–414, 2014.
- [SKZ17] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Spectral bounds for the ising ferromagnet on an arbitrary given graph. *Journal of Statistical Mechanics*, 2017(5), 2017.
- [SM07] Charles Sutton and Andrew McCallum. Improved dynamic schedules for belief propagation. In *Proceedings of UAI*, 2007.
- [SRF16] Christopher Srinivasa, Siamak Ravanbakhsh, and Brendan Frey. Survey propagation beyond constraint satisfaction problems. In *Proceedings of AISTATS*, pages 286–295, 2016.
- [Ste04] Hans J Stetter. *Numerical Polynomial Algebra*, volume 85. SIAM, 2004.
- [Str16] Gilbert Strang. *Introduction to Linear Algebra*. Wellesley-Cambridge Press, 5 edition, 2016.
- [SW05] Andrew J Sommese and Charles W Wampler. *The Numerical Solution of Systems of Polynomials Arising in Engineering and Science*, volume 99. World Scientific, 2005.
- [Tes12] Gerald Teschl. *Ordinary Differential Equations and Dynamical Systems*. American Mathematical Society, 2012.
- [TJ02] Sekhar C Tatikonda and Michael I Jordan. Loopy belief propagation and Gibbs measures. In *Proceedings of UAI*, pages 493–500, 2002.
- [TM06] Nobuyuki Taga and Shigeru Mase. Applications of Gibbs measure theory to loopy belief propagation algorithm. In *Mexican International Conference on Artificial Intelligence*, pages 197–207, 2006.

- [TR06] Peng Hui Tan and Lars K Rasmussen. Belief propagation for coded multiuser detection. In *IEEE International Symposium on Information Theory*, pages 1919–1923, 2006.
- [Ver99] Jan Verschelde. Algorithm 795: Phcpack: A general-purpose solver for polynomial systems by homotopy continuation. *ACM Transactions on Mathematical Software (TOMS)*, 25(2):251–276, 1999.
- [Vit67] Andrew Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269, 1967.
- [Vol17] Voltaire. *Oeuvres Complètes*. 1817.
- [Von13] Pascal O Vontobel. Counting in graph covers: A combinatorial characterization of the bethe entropy function. *IEEE Transactions on Information Theory*, 59(9):6018–6048, 2013.
- [Wai85] Tom Waits. *Rain Dogs*. Island Records, 1985.
- [Wai02] Martin Wainwright. *Stochastic processes on graphs with cycles: geometric and variational approaches*. PhD thesis, Massachusetts Institute of Technology, 2002.
- [Wat10] Yusuke Watanabe. *Discrete geometric analysis of message passing algorithm on graphs*. PhD thesis, The Graduate University for Advanced Studies, 2010.
- [WD16] Adrian Weller and Justin Domke. Clamping improves TRW and mean field approximations. In *Artificial Intelligence and Statistics*, pages 38–46, 2016.
- [Wei00] Yair Weiss. Correctness of local probability propagation in graphical models with loops. *Neural Computation*, 12(1), 2000.
- [Wel14] Adrian Weller. *Methods for Inference in Graphical Models*. PhD thesis, Columbia University, 2014.
- [Wel16] Adrian Weller. Uprooting and rerooting graphical models. In *Proceedings of ICML*, pages 21–29, 2016.
- [WF09] Yusuke Watanabe and Kenji Fukumizu. Graph zeta function in the Bethe free energy and loopy belief propagation. In *Proceedings of NIPS*, pages 2017–2025, 2009.
- [Wig58] Eugene P Wigner. On the distribution of the roots of certain symmetric matrices. *The Annals of Mathematics*, 67(2):325–327, 1958.
- [WJ08] Martin J Wainwright and Michael I Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.
- [WJ13] Adrian Weller and Tony Jebara. Bethe bounds and approximating the global optimum. In *Proceedings of AISTATS*, pages 618–631, 2013.
- [WJ14a] Adrian Weller and Tony Jebara. Approximating the Bethe partition function. In *Proceedings of UAI*, 2014.
- [WJ14b] Adrian Weller and Tony Jebara. Clamping variables and approximate inference. In *Proceedings of NIPS*, pages 909–917, 2014.

- [WJW02] Martin J Wainwright, Tommi Jaakkola, and Alan S Willsky. Tree-based reparameterization for approximate inference on loopy graphs. In *Proceedings of NIPS*, pages 1001–1008, 2002.
- [WJW03a] Martin J Wainwright, Tommi S Jaakkola, and Alan S Willsky. Tree-based reparameterization framework for analysis of sum-product and related algorithms. *IEEE Transactions on information theory*, 49(5):1120–1146, 2003.
- [WJW03b] Martin J Wainwright, Tommi S Jaakkola, and Alan S Willsky. Tree-reweighted belief propagation algorithms and approximate ML estimation by pseudo-moment matching. In *Proceedings of AISTATS*, 2003.
- [WJW05] Martin J Wainwright, Tommi S Jaakkola, and Alan S Willsky. A new class of upper bounds on the log partition function. *IEEE Transactions on Information Theory*, 51(7):2313–2335, 2005.
- [WSW08] Alan S Willsky, Erik B Sudderth, and Martin J Wainwright. Loop series and Bethe variational bounds in attractive graphical models. In *Proceedings of NIPS*, pages 1425–1432, 2008.
- [WT01] Max Welling and Yee Whye Teh. Belief optimization for binary networks: A stable alternative to loopy belief propagation. In *Proceedings of UAI*, pages 554–561, 2001.
- [WT03] Max Welling and Yee Whye Teh. Approximate inference in Boltzmann machines. *Artificial Intelligence*, 143(1):19–50, 2003.
- [WTJS14] Adrian Weller, Kui Tang, Tony Jebara, and David Sontag. Understanding the Bethe approximation: When and how can it go wrong? In *Proceedings of UAI*, pages 868–877, 2014.
- [Wym07] Henk Wymeersch. *Iterative Receiver Design*, volume 234. Cambridge University Press Cambridge, 2007.
- [YFW01] Jonathan S Yedidia, William T Freeman, and Yair Weiss. Generalized belief propagation. In *Proceedings of NIPS*, pages 689–695, 2001.
- [YFW05] Jonathan S Yedidia, William T Freeman, and Yair Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51(7):2282–2312, 2005.
- [You98] Allan Peter Young. *Spin Glasses and Random Fields*, volume 12. World Scientific, 1998.
- [YR03] Alan L Yuille and Anand Rangarajan. The concave-convex procedure. *Neural Computation*, 15(4):915–936, 2003.
- [ZK16] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.