Correlated Equilibria for Approximate Variational Inference in MRFs

Luis E. Ortiz LEORTIZ@UMICH.EDU

Boshen Wang BOSHENW@UMICH.EDU

University of Michigan - Dearborn

Ze Gong ZGONG11@ASU.EDU

Arizona State University

Abstract

Almost all of the work in graphical models for game theory has mirrored previous work in probabilistic graphical models. Our work considers the opposite direction: Taking advantage of advances in equilibrium computation for probabilistic inference. In particular, we present formulations of inference problems in Markov random fields (MRFs) as computation of equilibria in a certain class of game-theoretic graphical models. While some previous work explores this direction, we still lack a more precise connection between variational probabilistic inference in MRFs and correlated equilibria. This paper sharpens the connection, which helps us exploit relatively more recent theoretical and empirical results from the literature on algorithmic and computational game theory on the tractable, polynomial-time computation of exact or approximate correlated equilibria in graphical games with arbitrary, loopy graph structure. Our work discusses how to design new algorithms with equally tractable guarantees for the computation of approximate variational inference in MRFs. In addition, inspired by a previously stated game-theoretic view of tree-reweighted message-passing techniques for belief inference as a zero-sum game, we propose a different, general-sum potential game to design approximate fictitious-play techniques. Empirical evaluations on synthetic experiments and on an application to soft de-noising on real-world image datasets illustrate the performance of our proposed approach and shed some light on the conditions under which the resulting belief inference algorithms may be most effective relative to standard state-of-the-art methods.

Keywords: game theory; belief inference; correlated equilibria; equilibrium computation

1. Introduction

Almost all of the work in graphical games has borrowed heavily from analogies to *probabilistic graphical models (PGMs)*. However, over-reliance on those analogies and previous standard approaches to exact inference might have led that approach to face the same computational roadblocks that plagued most exact-inference techniques.

As an example of work that heavily exploits previous work in PGMs, Kakade et al. (2003) designed polynomial-time algorithms based on linear programming for computing CE in standard graphical games with tree graphs. The approach and polynomial-time results extend to graphical games with bounded-tree-width graphs and graphical polymatrix games with tree graphs. Exact inference is tractable in PGMs whose graphs have bounded treewidth, but intractable in general. Papadimitriou and Roughgarden (2008) showed the intractability of computing the "social-welfare" optimum *correlated equilibria (CE)* in arbitrary graphical games. Everything seemed to point toward an eventual resignation that the approach of Kakade et al. (2003), along with any other approach to the problem for that matter, had hit the "bounded-treewidth-threshold wall."

Yet, soon after, Papadimitriou (2005) took a radically different approach to the problem, and surprised the community with an efficient algorithm for computing CE not only in graphical games,

but also in almost all known compactly representable games. Jiang and Leyton-Brown (2015) built upon Papadimitriou's idea to provide an improved polynomial-time algorithm.

An immediate question that arises from the algorithmic results just described is, what is so fundamentally different between the problem of exact inference in graphical models and equilibrium computation that made this result possible in the context of graphical games? It is well-known that *pure strategy Nash equilibrium (PSNE)* is inherently a classical/standard discrete *constraints satisfaction problem (CSP)*. It is also well-known that any CSP with a solution can be cast as a most-likely, or equivalently, a *maximum a posteriori (MAP)* assignment estimation problem in *Markov random fields (MRFs)*. Through this connection, it is clear that there exists a MAP formulation of PSNE. But what about other, more general forms of equilibria?

We present here a formulation of the problem of equilibrium computation as a kind of local conditions for different approximations to belief inference. Similarly, we show how one can view some special games, called *graphical potential games* (Ortiz, 2015), as defining an *equivalent* MRF whose "locally optimal" solutions correspond to *arbitrary* equilibria of the game. We employ ideas from the literature of learning in games (Fudenberg and Levine, 1999), particularly no-regret algorithms and fictitious play, to propose two specific instances of game-theoretic inspired, practical, and effective heuristics for belief inference in MRFs, based on two different approaches: one local, the other global. We evaluate our proposed algorithms within the context of the most popular, standard, and state-of-art techniques from the literature in PGMs.

2. Preliminaries

Denote by $x \equiv (x_1, x_2, \ldots, x_n)$ an n-dimensional vector and by $x_{-i} \equiv (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$ the same vector without component i. Similarly, for every set $S \subset [n] \equiv \{1, \ldots, n\}$, denote by $x_S \equiv (x_i : i \in S)$ the (sub-)vector formed from x using only components in S, such that, letting $S^c \equiv [n] - S$ denote the complement of S, we can denote $x \equiv (x_S, x_{S^c}) \equiv (x_i, x_{-i})$ for every i. If $A_1, \ldots, A_n$ are sets, denote by $A \equiv \times_{i \in [n]} A_i$, $A_{-i} \equiv \times_{j \in [n]-\{i\}} A_j$ and $A_S \equiv \times_{j \in S} A_j$.

Let $G = (V, E)$ be an undirected graph, with finite set of n *vertices* or *nodes* $V = \{1, \ldots, n\}$ and a set of (undirected) edges E. For each node i, let $\mathcal{N}(i) \equiv \{j \mid (i, j) \in E\}$ be the set of neighbors of i in G, *not including* i, and $N(i) \equiv \mathcal{N}(i) \cup \{i\}$ the set *including* i. A *clique* C of G is a set of nodes with the property that they are all mutually connected: for all $i, j \in C$, $(i, j) \in E$; in addition, C is *maximal* if there is no other node k outside C that is also connected to each node in C: for all $k \in V - C$, $(k, i) \notin E$ for some $i \in C$.

Denote by Ω_i the (finite) set of possible values that random variable X_i can take. By definiton, a joint probability distribution P over Ω is a *Markov random field (MRF) with respect to (wrt) an undirected graph* G if for all $x \in \Omega$, for every node i, $P(X_i = x_i \mid X_{-i} = x_{-i}) = P(X_i = x_i \mid X_{\mathcal{N}(i)} = x_{\mathcal{N}(i)})$. In that case, the neighbors/variables $X_{\mathcal{N}(i)}$ form the *Markov blanket* of node/variable X_i. Define *local potential functions* $\{\phi_C : \Omega_C \to \mathbb{R}\}$ over each clique C in the graph, and $\Psi(x) \equiv \sum_{C \in \mathcal{C}} \phi_C(x_C)$ as a *Gibbs potential* with respect to G. By the Hammersley and Clifford (1971)'s Theorem, we can express a positive MRF as $P(X = x) \propto \exp(\Psi(x))$.

Game theory (von Neumann and Morgenstern, 1947) provides a mathematical model of the stable behavior (or outcome) that may result from the interaction of rational individuals. This paper concentrates on *noncooperative* settings: individuals maximize their *own* utility, act *independently*, and do not have (direct) control over the behavior of others. Let $V = [n]$ denote a finite set of n players in a game. For each player $i \in V$, let A_i denote the set of *actions* or *pure strategies* that i

can play. The set A is the set of *joint actions*. Let $x \in A$ denote a *joint action*, and x_i the *individual action* of player i in x. Let $M_i : A \to \mathbb{R}$ denote the *payoff/utility function* of player i. If the A_i's are finite, then M_i is called the *payoff matrix* of player i of a game in *normal-form*.

Graphical potential games (GPGs), of which Ortiz (2015) provides a characterization, play a key role in establishing a stronger connection between probabilistic inference in MRFs and equilibria in games than previously noted. While not all games in normal-form have a PSNE (e.g., rock-paper-scissors), GPGs always have a PSNE.

Equilibria are generally considered *the* solutions of games. Various notions of equilibria exist. Given $\epsilon \geq 0$, an *ϵ-approximate correlated equilibrium (ϵ-CE)* is a joint probability distribution Q over A such that, for all players i, $x_i, x_i' \in A_i$, $x_i \neq x_i'$, and $Q(x_i) > 0$, $\sum_{x_{-i}} Q(x_1, x_{-i}) M_i(x_i, x_{-i}) \geq \sum_{x_{-i}} Q(x_i', x_{-i}) M_i(x_i', x_{-i}) - \epsilon$, where $Q(x_i) \equiv \sum_{x_{-i}} Q(x_i, x_{-i})$ is the (marginal) probability that player i will play x_i according to Q. An *ϵ-approximate mixed strategy Nash equilibrium (ϵ-MSNE)* is and ϵ-CE that is a *product* distribution. An *ϵ-approximate pure strategy (Nash) equilibrium (PSNE)* is a joint action x^* for which $Q(x) = \mathbb{1}[x \in x^*]$ is a an ϵ-CE. An *exact* CE (Aumann, 1974), MSNE, and PSNE are a 0-CE, 0-MSNE, and 0-PSNE, respectively. *Every* game in normal-form has at least one MSNE (Nash, 1950, 1951). Thus, every game has an ϵ-MSNE "solution," and thus also an ϵ-CE.

3. Equilibria and Inference

In this section, we seek to provide a more formal connection between probabilistic and game-theoretic graphical models in order to leverage advances in equilibrium computation for probabilistic inference. (We omit all proofs and derivations due to space constraints.)

3.1 PSNE and Approximate MAP Inference

Consider an MRF P with respect to graph G and Gibbs potential Ψ defined by the set of potential functions $\{\phi_C\}$. For each node i, denote by $\mathcal{C}_i \subset \mathcal{C}$ the subset of cliques in G that include i. Note that the (inclusive) neighborhood of player is given by $N(i) = \cup_{C \in \mathcal{C}_i} C$. Define an *MRF-induced GPG* with the same graph G, and for each player i, hypergraph with hyperedges $\mathcal{C}_i$ and local-clique payoff hypermatrices $M'_{i,C}(x_C) \equiv \phi_C(x_C)$ for all $C \in \mathcal{C}_i$ defining the Gibbs potential function Ψ: $M_i(x_i, x_{-i}) - M_i(x_i', x_{-i}) = \Psi(x_i, x_{-i}) - \Psi(x_i', x_{-i}) = \sum_{C \in \mathcal{C}_i} M'_{i,C}(x_i, x_{C-\{i\}}) - M'_{i,C}(x_i', x_{C-\{i\}})$.

Property 1 *The representation size of the MRF-induced game is the same as that of the MRF.*

Similarly, for any potential game, one can define a *game-induced MRF* using the potential function of the game whose set of local maxima corresponds exactly to the set of PSNE of the potential game. Through this connection we can show that solving the local-MAP problem in MRFs is PLS-complete in general (Fabrikant et al., 2004). In short, we can use algorithms for PSNE as heuristics to compute locally optimal MAP assignments of P and *vice versa*.

3.2 CE and Belief Inference

Going beyond PSNE and MAP estimation, this subsection begins to establish a stronger, and potentially more useful connection between probabilistic inference and more general concepts of equilibria in games. For any pair of distributions P and Q, let $H(Q, P) \equiv \sum_x Q'(x)[-\ln P(x)]$ and

$\mathrm{KL}(Q \parallel P) \equiv \sum_x Q(x) \ln(Q(x)/P(x)) = H(Q,P) - H(Q)$ be the *cross entropy* and *Kullback-Leibler divergence* between probability distributions Q and P, with respect to Q, respectively.

Property 2 *For any MRF P, any CE Q of the game induced by P satisfies* $\mathrm{KL}(Q \parallel P) = \min_{Q'_{X'_i|X_i}} \mathrm{KL}(\widetilde{Q}^i \parallel P) + H(\widetilde{Q}^i) - H(Q)$ *for all i, with* $\widetilde{Q}^i(x) \equiv \sum_{x'_i} Q'_{X'_i|X_i}(x_i|x'_i) Q(x_{-i}|x'_i) Q(x'_i)$.

Hence, *any CE of the MRF-induced game is an approximate local optimum (or critical point) of an approximation of the MRF based on variational inference, where the quality of the approximate local optimum depends on the resulting difference in entropies $H(\widetilde{Q}^i) - H(Q)$ at the local optimum.*

We can also derive a slightly weaker condition based on a relaxation of CE called a *coarse CE* (Hannan, 1957; Moulin and Vial, 1978). For each player i, denote by $Q_{-i}(x_{-i}) \equiv \sum_{x_i} Q(x_i, x_{-i})$ the marginal distribution of play over the joint-actions of all players *except* player i; by $Q'_i Q_{-i}$ the joint distribution defined as $(Q'_i Q_{-i})(x) \equiv Q'_i(x_i) Q_{-i}(x_{-i})$ for all x; and by $H(Q_{i|-i}) \equiv H(Q_{-i}) - H(Q)$ the conditional entropy of the individual play of player i given the joint play of all the players except i.

Property 3 *For any MRF P, any CE Q of the game induced by P satisfies* $\mathrm{KL}(Q \parallel P) \leq \min_i \left[\min_{Q'_i} \mathrm{KL}(Q'_i Q_{-i} \parallel P) + H(Q'_i) \right] - H(Q_{i|-i})$.

This connection allows us to apply several simple methods for computing coarse CE.

In the special case of MSNE, the joint mixed strategy $Q(x) = \prod_i Q_i(x_i)$ is a product distribution and we can obtain a tighter condition than for arbitrary (coarse) CE. Denote by $Q^{\times}_{-i}(x_{-i}) \equiv \prod_{j \neq i} Q_j(x_j) = \sum_{x_i} Q(x)$ the (marginal) joint action of play over all the players except i, and denote by $(Q'_i Q^{\times}_{-i})$ the probability distribution defined such that the probability of x is $(Q'_i Q^{\times}_{-i})(x) \equiv Q'_i(x_i) Q^{\times}_{-i}(x_{-i})$.

Property 4 *For any MRF P, any MSNE Q of the game induced by P satisfies* $\mathrm{KL}(Q \parallel P) = \left[\min_{Q'_i} \mathrm{KL}(Q'_i Q^{\times}_{-i} \parallel P) + H(Q'_i) \right] - H(Q_i)$*, for all i.*

Hence, *an MSNE Q of the game is an approximate locally-optimal mean-field approximation*. In fact, based on the last property, we can show that Q is a so-called *logit equilibrium*, which is an instance of *quantal response equilibria (QRE)* (McKelvey and Palfrey, 1995), and can equivalently be viewed as a PSNE of a so-called *infinite* potential game in this case, if and only if it is a local optimum (or critical point) of a mean-field approximation. This establishes connections to the literature on learning in games (Fudenberg and Levine, 1999), which opens up opportunities to use algorithms from that community, such as those based on fictitious play, to perform probabilistic inference, as we do here. The last property also suggets that we could use appropriately-modified versions of algorithms for MSNE as heuristics to obtain a mean-field approximation of the true marginals. Going in the opposite direction, by treating any (graphical) potential game as an MRF, for any fixed $\lambda > 0$, logistic fictitious play in any potential game converges to an approximate $(\lambda / \min_i |A_i|)$-MSNE of the potential game. There has been recent work in this direction, which explores the connection between learning in games and mean-field approximations in machine learning (Rezek et al., 2008). That work establishes the same connection to logit equilibrium and proposes new algorithms based on fictitious play for simple mean-field approximation applied to statistical (Bayesian) estimation.

We now return to the connection between arbitrary CE and other higher-order variational approximations and discuss some computational implications. Besides Kakade et al. (2003), Ortiz

et al. (2007) proposed the principle of maximum entropy (MaxEnt) for equilibrium selection of CE in graphical games. They studied several properties of the MaxEnt CE, designed a monotonically increasing algorithm to compute it, and discussed a learning-dynamics view of the algorithm. Kamisetty et al. (2011) employed advances in approximate inference methods to propose approximation algorithms to compute CE. In all of those cases, the general approach is to use ideas from PGMs to design algorithms to compute CE. The focus of this paper is the opposite direction: employing ideas from game theory to design algorithms for belief inference in PGMs.

Consider the algorithms of Papadimitriou (2005) or Jiang and Leyton-Brown (2015) (see also Papadimitriou and Roughgarden (2008) and Jiang and Leyton-Brown (2011)), which we can use to compute a CE of the MRF-induced game in polynomial time. Such CE will be, by construction, also a *(polynomially-sized) mixture of product distributions*. (In the case of Jiang and Leyton-Brown's algorithm it will be a mixture of a subset of the joint-action space, which is equivalent to a probability mass function over a *polynomially-sized* subset of the joint-action space; said differently, a mixture of product of indicator functions, each product corresponding to particular outcomes of the joint-action space.) Hence, both algorithms provide a means to obtain a heuristic estimate of a local optimum (or critical point) of such a mixture *in polynomial time.*

Finally, this connection also suggests that we can (in principle) use any learning algorithm that guarantees convergence to the set of CE as a heuristic for approximate inference. Several so-called "no-regret" learning algorithms satisfy those conditions. Indeed, we use a variant of a simple version of a "no-regret" algorithm in our experiments. Because only approximate equilibria is typically guaranteed in finite time in practice, there may be an additional "gap" in the local conditions of each player i given in Properties 2, 3, and 4 that equals the player's "regret" $\epsilon_i \equiv \max_{x_i'} \sum_x Q(x) \left(\Psi(x_i', x_{-i}) - \Psi(x)\right)$, where the precise expression for Q depends on the particular property used. We do not discuss this further here but briefly note that in practice "regret gaps" seem to have a larger effect than "entropy-difference gaps" and that the effect seems algorithm-dependent.

3.3 Approximate Fictitious Play in a Two-Player Potential Game for Belief Inference

This section presents a game-theoretic fictitious-play approach to estimation of node-marginal probabilities in MRFs. The approach this time is more global in terms of how we use the whole joint-distribution for the estimation of individual marginal probabilities. The inspiration for the approach presented here follows from the work of Wainwright et al. (2005). The section concentrates on Ising models, an important, special MRF instance from statistical physics with its own interesting history.

Definition 1 *An* Ising model (IM) *wrt an undirected graph $G = (V, E)$ is an MRF wrt G such that $P_\theta(x) \propto \exp\left(\sum_{i \in V} b_i x_i + \sum_{(i,j) \in E} w_{i,j} x_i x_j\right)$ where $\theta \equiv (\mathbf{b}, \mathbf{W})$ is the set of node biases b_i's and edge-weights w_{ij}'s, which are the parameters defining the joint distribution P_θ over $\{-1, +1\}^n$.*

Denote by $\mathbb{T}_G$ the set of all spanning trees of connected (undirected) graph $G = (V, E)$ that are maximal with respect to E (i.e., does not contain any spanning forests). If spanning tree $T \in \mathbb{T}_G$, we denote by $E(T) \subset E$ the set of edges of T.

Initialize $x^{(1)} \leftarrow \text{Uniform}(\{-1, +1\}^n)$, and for each $(i, j) \in E$, $\widehat{\mu}^{(1)}_{(i,j)} \leftarrow x^{(1)}_i x^{(1)}_j$. At each iteration $l = 1, 2, \ldots, m$,

1: $\mathbb{T}^* \leftarrow \arg\max_{T \in \mathbb{T}_G} \sum_{(i,j) \in E} \mathbb{1}[(i, j) \in E(T)]\, w_{ij} \widehat{\mu}^{(l)}_{(i,j)}$

2: $T^{(l)} \leftarrow \text{Uniform}\,(\mathbb{T}^*)$

3: $s_l \leftarrow \text{Uniform}(\{1,\ldots,l\})$
4: $\mathcal{X}^* \leftarrow \arg\max_{x\in\{-1,+1\}^n} \sum_{i\in V} b_i x_i + \sum_{(i,j)\in E(T^{(s_l)})} w_{ij} x_i x_j$
5: $x^{(l+1)} \leftarrow \text{Uniform}\,(\mathcal{X}^*)$
6: **for all** $(i,j) \in E$ **do**
7: $\quad v_{(i,j)}^{(l+1)} \leftarrow x_i^{(l+1)} x_j^{(l+1)} \times \begin{cases} 1, & \text{if MSNE,} \\ \mathbb{1}\big[(i,j) \in E(T^{(s_l)})\big], & \text{if CE} \end{cases}$
8: $\quad \widehat{\mu}_{(i,j)}^{(l+1)} \leftarrow \frac{l\,\widehat{\mu}_{(i,j)}^{(l)} + v_{(i,j)}^{(l+1)}}{l+1}$
9: **end for**

For each IM's random-variable index $i = 1,\ldots,n$, set $p_i^{(m+1)} = \frac{1}{m+1}\sum_{l=1}^{m+1} \mathbb{1}\left[x_i^{(l)} = 1\right]$ as the estimate of the exact IM's marginal probability $p_i \equiv P_\theta(X_i = 1)$.

Within the literature on PGMs, Hamze and de Freitas (2004) propose an MCMC approach based on sampling non-overlapping trees. While our approach has a sampling flavor, its exact connection to MCMC is unclear at best. Also, the spanning trees that our algorithm generates may overlap.

The game implicit in the heuristic algorithm above is a two-player potential game between a *"joint-assignment" (JA)* player and a *"spanning-tree" (ST)* player. The potential function is $\Psi_{X,\mathcal{T}}(x,T) = \sum_{i\in V} b_i x_i + \sum_{(i,j)\in E} \mathbb{1}[(i,j) \in E(T)] w_{ij} x_i x_j$. The payoff functions M_X and $M_\mathcal{T}$ of the JA player and the ST player, respectively, are identical and equal the potential function $\Psi_{X,\mathcal{T}}(x,T)$: formally, $M_X(x,T) = M_\mathcal{T}(x,T) = \Psi_{X,\mathcal{T}}(x,T)$. Note that the payoff function of the ST player is *strategically equivalent* to the function $\sum_{(i,j)\in E} \mathbb{1}[(i,j) \in E(T)] w_{ij} x_i x_j$.

Technically, this is a game with identical payoffs, which are known to have what Monderer and Shapley (1996) called the *fictitious play property*: the empirical play of fictitious play is guaranteed to converge to an MSNE of the game. In fictitious play, each player uses the *empirical distribution of play* as an estimate or belief of how the other player would behave in the future, *not just* the other player's *last action* as in sequential best-reply. Determining a best-response for the ST player is easy (i.e., using an algorithm for computing maximal spanning tree such as Kruskal's). Unfortunately the same is in general not possible for the JA player, whose best-response is as hard as computing a MAP assignment of another IM with the same graph and (generally non-zero) bias/node parameters, but a slightly different set of edge-weights. That is why we draw one tree uniformly at random from the empirical distribution and find a best-response to that tree as a best-response from the JA player. Such an approach is equivalent to a type of smooth best-response. If both players were to do the same, *simultaneously*, the result is a stochastic version of fictitious play or *stochastic fictitious play* for short (Fudenberg and Levine, 1999). The empirical distribution of play of stochastic fictitious play in a game with identical payoffs, or what's strategically equivalent, any potential game, also converges to an MSNE of the game (Hofbauer and Sandholm, 2002). In our case, however, we really have a type of "hybrid" sequential-version, where the ST player is always behaving as in standard fictitious play, while the JA player is behaving according to a stochastic fictitious play.

Within the context of potential games, while *sequential best-reply* converges to a PSNE (i.e., a joint assignment), *fictitious play* can converge to an MSNE of the game. Monderer and Shapley (1996) provide an example in a 2-player 2-action normal-form (coordination) game with identical payoffs. Said differently, the resulting empirical distribution of play for the JA player may be to what Monderer and Shapley (1996) themselves call a "purely mixed strategy" (i.e, every action is played with positive probability).

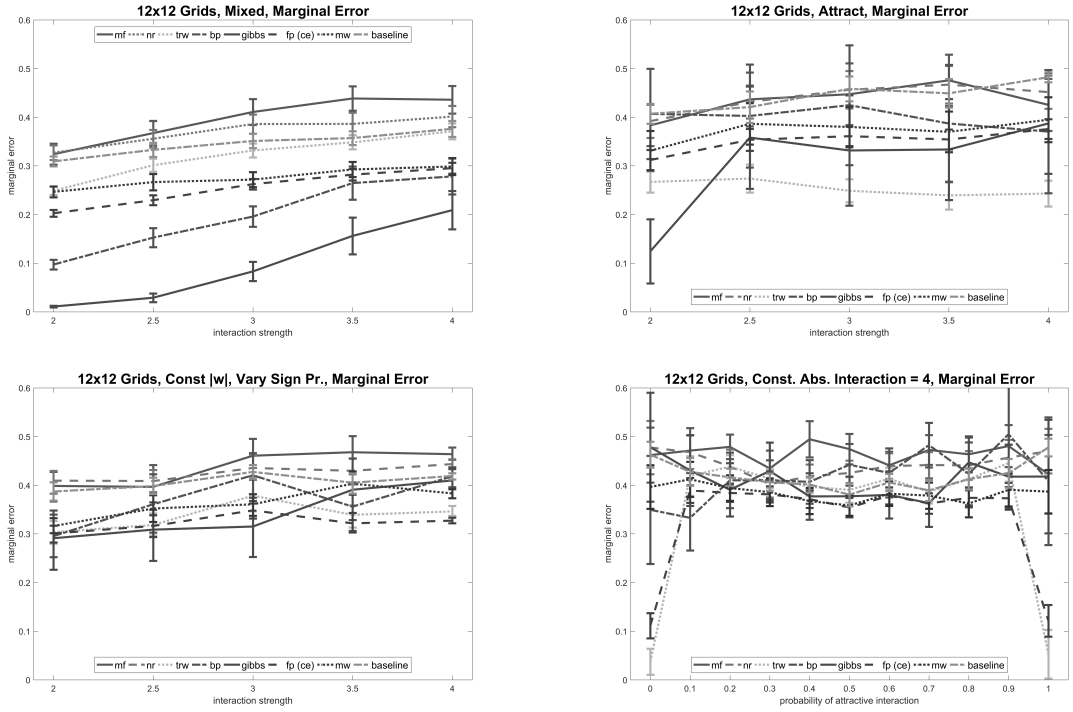

Figure 1: **Evaluation on IMs with 12×12 Grids.** The top plots are for the so-called "mixed" (*Left*) and "attractive" (*Right*) instances of IMs.For both plots, the x-axis is the largest magnitude of the edge-weights: i.e., $w = \max_{(i,j)\in E} |w_{ij}|$. The y-axis is the average, over 50 randomly generated IMs, of the *average*, over all of the 144 variables, of the absolute difference between the estimate and exact marginal probability for the random variable corresponding to that node. The bottom plots are from an alt. evaluation of IMs. (*Bottom-left plot*) Here, the edge-weight *magnitude* w is constant for each interaction strength, and nr uses 10^6 iterations. For all cases, the result is the average over all values of the probability of attractive interaction q, and over 5 IMs for each q; *except* for the case of constant edge-weight magnitude $w = 4$, which uses 50 IMs for each q. (*Bottom-right plot*) Results for each q value with $w = 4$, with 50 IMs as samples for each.

4. Experiments

In this section we present the results of experiments on the performance of the game-theoretic-inspired heuristics we propose in this paper for approximate belief inference in MRFs. Our algorithms have very simple implementations. We also compare them with the most popular approximation algorithms and heuristics, with equally simple implementations.

4.1 Experiments on Randomly-Generated Synthetic Models

The experimental design in terms of the class of IMs is as in Domke and Liu (2013). We consider IMs with $d \times d$ simple grid graphs, which are planar (i.e., $n = d^2$). For each value $w \in \{2, 2.5, 3, 3.5, 4\}$, we generated random IMs with edge-weights $w_{ij} \sim \text{Uniform}([-w, w])$ or $w_{ij} \sim \text{Uniform}([0, w])$ for the "mixed" or "attractive" case, respectively, for each $(i, j) \in E$, i.i.d., and node biases $b_i \sim \text{Uniform}([-1, +1])$, also i.i.d. for all i, and independent of the edge-weights. We

also considered for evaluation a class of IMs in which edge-weights have constant *magnitude* (i.e., $w = \max_{(i,j)\in E} |w_{ij}|$), but in which we vary the probability q of attractive edge-weights; that is, given a probability q, the sign of the edge-weight are i.i.d. random variables in which the sign is positive with probability q. For evaluation using this class, we consider $w \in \{2, 2.5, 3, 3.5, 4\}$. For each q, we randomly generated 50 IMs as samples for $w = 4$, and 5 samples for each $w \neq 4$.

We evaluated several different so-called "no-regret" algorithms for learning in games, with varying degrees of equilibrium-convergence guarantees, as approximation for belief inference. Each iteration of these algorithms takes roughly the same amount of time as that for Gibbs sampling. One set of such algorithms we used is based on the *multiplicative weight update (MWU)* algorithm (Blum and Mansour, 2007). In our implementation of the MWU algorithm, for each player i at each round $t \geq 1$, we set the probability of playing action x_i at round $t+1$, which we denote by $x_i^{(t+1)}$, to be $p_{t+1}(x_i) \propto p_t(x_i)\left(1 - \eta_t(1 - \bar{M}_i(x_i, x_{-i}^{(t)}))\right)$, where $\eta_t = \sqrt{\frac{\ln(2)}{t}}$ is analogous to a learning rate in ML, and $\bar{M}_i$ is the normalized payoff function for player i. In this case, the players have no external regret, so we label the algorithm as "mw_er". We also run a variant of mw_er with a constant $\eta_t = 0.01$, labeled as "mw_er_cf," in which players have approximately no external regret. The minimization of external regret means mw_er has guaranteed convergence to the set of coarse CE, whereas mw_er_cf converges to the set of approximate coarse CE. MWU can also be adapted to minimize swap regret instead of external regret. In that case, it would converge to the set of CE or approximate CE, depending on if η_t is a constant or not. We include "mw_sr" and "mw_sr_cf" implementations, the former using $\eta_t = \sqrt{\frac{\ln(2)}{t}}$ as above and the latter using a constant $\eta_t = 0.01$.

We also consider a simple version of an (approximate) no-swap-regret algorithm by Hart and Mas-Colell (2000), which we denote as "nr" from now on. We set the number of iterations of the nr algorithm to 10^5 for the standard experimental setup, and to 10^6 for our proposed new evaluation setting. Also, we only present results for the sequential, "semi-stochastic" fictitious play we propose for the 2-player potential game for the case of CE only, which we denote as "fp (ce)" from now on. We set the number of iterations $m = 15$. (We omit the results for the MSNE-instantiation of that fictitious play algorithm because they were similar to those for fp (ce), at least for $m = 15$.)

We compared the different mw-type algorithms, the simple nr algorithm, and our proposed fp (ce) to (1) standard mean-field approximation (mf), with sequential/axis-parallel updates; (2) standard belief propagation (bp), with simultaneous updates; (3) TRW (trw); and (4) the Gibbs sampler (gs). As a baseline (bl), we use the simplest possible estimator from the perspective of average marginal-error to measure quality: always use 0.5 as the estimate of the exact marginal distribution of each variable. Among mw-type algorithms, we only present the results of mw_er_cf, referred to simply as "mw" from now on, because it outperforms all other types almost consistently.

We evaluate mean field (mf) using sequential/axis-parallel updates, stopping if the max. abs. diff. in prob. values between iters. is $\leq 10^{-5}$, and using a max. num. of iters. $= 10^6$. For belief propagation (bp) we use simultaneous updates, and "smooth" the update based on the avg. of the current value and the new value in order to "dampen" or at least try to prevent oscillations and improve the likelihood of convergence, stopping if the max. abs. diff. in prob. values between iters. is $\leq 10^{-7}$, and a max. num. of iters. $= 10^5$. For tree reweighted message-passing (trw), we use a constant parameter $\rho = 0.55$ for all corresponding edge-appearance-prob. parameters ρ_{ij}'s (Wainwright et al., 2005), along with a smooth update and the same stopping criterion as for bp. We use 10^6 iters. for the Gibbs sampler.

The top plots in Fig. 1 summarize our results for the most common classes of IMs considered in the experimental evaluation of approximation algorithms and heuristic for belief inference in the literature as described above. We perform hypothesis testing for the result in these classes of IMs using paired z-tests on the individual (i.e., not joint) differences, each with p-value 0.05. Hence, all the statements are statistically significant with respect to such hypothesis tests. Note that there is no overall best approximation technique for these classes.

"Mixed" Case (Top-Left Plot, Fig. 1). gs is best for all w in this case. Among the other approximation algorithms, we observe the following: (1) fp (ce) is worse than bp for $w < 3.5$, and indistinguishable from bp for $w \geq 3.5$; (2) fp (ce) is consistently better than trw; (3) trw is consistently worse than bp; (4) all methods, except for mf and nr, are consistently better than bl; mf and nr are consistently worse than bl; (5) mf is better than mr for $w \geq 3$, and indistinguishable from each other for $w < 3$; (6) mw is worse than fp (ce) for $w < 3$, but indistinguishable from fp (ce) for $w \geq 3$.

"Attractive" case (Top-Right Plot, Fig. 1). There is no clear overall best. We also observe the following: (1) trw is best among all methods for $w \geq 3.0$, indistinguishable from gs for $w = 2.5$, and worse than gs for $w = 2.0$; (2) fp (ce) is worse than gs for $w = 2.0$, but better than gs for $w = 4.0$, and indistinguishable from gs otherwise; (3) mf, nr, and bp are consistently indistinguishable from each other, except for $w = 2$ where bp is better than nr; (4) bp and bl are consistently indistinguishable, except for $w = 4.0$, where bp is better; (5) mw and fp (ce) are consistently indistinguishable.

The bottom plots in Fig. 1 summarize our experimental results for a class of IMs which appears to lead to "harder" instances. Note that the standard 95% CIs based on a Gaussian approximation resulting from CLT do not directly apply here because the averages are over different q values, each of which may have different distributional properties (e.g., different variances). For $w < 4$, because we are computing the average marginal-error over every q, each based on only 5 samples, we use the bootstrap method to compute the 95% CIs over the overall average for each method and each w, using 100 samples. For $w = 4$, because we have 50 samples for each q, we use a properly adapted version of the standard 95% CIs which modifies the calculation of the overall variance to account for distributional differences from each q. We perform hypothesis testing for the result in these classes of IMs using two approaches depending on w. For $w = 4$, where we draw 50 models as samples for each q, we use appropriately modified paired z-tests on the individual (i.e., not joint) differences, each with p-value 0.05. We modify the calculation of the variances resulting from the average over the samples computed for each q. For $w < 4$, where we only draw 5 models as samples for each q, we use bootstrapped-based, individual, paired hypothesis-testing over each pair of aggregate differences between the methods for each of those values of w; we use 100 bootstrap samples, and p-value 0.05. All the statements are statistically significant with respect to such hypothesis tests.

Aggregate Results (Bottom-Left Plot, Fig. 1). There is no clear overall best. We also observe the following: (1) fp (ce) is best among all methods except for when $w = 2.0$, where gs is better; (2) trw is second best among all methods, except for when $w = 2.0$, where it is third best (behind fp (ce) and gs); (3) bp is consistently better than mf and nr except when $w = 3.5$, where it is indistinguishable from nr (but still better than mf); (4) mf is consistently worse than bl, except when $w = 4.0$, where they are indistinguishable. nr is also consistently worse than bl, except when $w = 2.5$, where they are indistinguishable; (5) gs is consistently better than mf, nr, and bl, except when $w = 4.0$, where gs and bl are indistinguishable.

Results for Const. Magnitude $w = 4$ (Bottom-Right Plot, Fig. 1). The results suggest that in fact such instances of IMs tend to be harder in the sense that even state-of-the-art algorithms such as TRW are no better than the simple baseline estimation, in which $\widehat{p}_i = 0.5$ for all nodes/variables i,

for all of the non-extreme values of q. In fact, the performance of TRW is *almost exactly the same* as baseline across the range of non-extreme values of q. (Note how the plot of the values for trw and bl are essentially on top of each other for values of q other than 0 or 1.) On the other hand, note how fp (ce) is consistently better than bl across the whole range of values for q. In fact, fp (ce) is always in the set of (statistically) best performers for all q with the exception of $q = 0.0$, where trw is better. Almost all the methods other than fp (ce) are no better, and often worse, than bl. The exceptions are that trw beats bl when $p \in \{0.0, 1.0\}$, bp beats bl when $p \in \{0.0, 0.2\}$, and gs beats bl when $p = 0.6$.

4.2 Experiments on More Realistic Models for (Soft) De-noising of BW Images

We use images of handwritten digits from the popular MNIST dataset (LeCun et al., 1998) to build IMs for soft de-noising: using the individual marginal probabilities as a confidence measure of the individual pixel values of the de-noised image. We only do qualitative evaluations here because of the intractability of exact inference in this case. The results shed some light on the type of inference problems for which some of the approximation methods proposed here may be most effective.

The MNIST images are 28×28 grayscale pixel images of handwritten digits. Therefore, we construct 28×28 simple planar IMs, again consisting of node biases b_i and edge weights w_{ij}. First, we convert the grayscale pixels to black (-1) and white (+1) using a threshold of 0.5. Next, for each pixel, we find the average value of that pixel across all training images, which is used as the node bias. Finally, we compute the average product between neighboring pixels, taken across all training images, *regardless of digit label*, and use it as the edge weight between those neighboring pixels.

We considered randomly-selected images from the MNIST test dataset, and convert them to black/white images using the same thresholding method as above. Then, we add noise to the black/white image by independently flipping each pixel value with some probability p (the "noise level"). Finally, we modify our existing node biases for each noisy image by taking into account our "observation" of the noisy image: If I denotes the matrix representation of the noisy black/white test image, the Gibbs potential of that image's posterior IM is $\Psi_I(x) \equiv \sum_{((i,j),(r,s))\in E} w_{(i,j),(r,s)} x_{(i,j)} x_{(r,s)} + \sum (i,j) \widetilde{b}_{(i,j)}(I(i,j))\, x_{(i,j)}$ where $\widetilde{b}_{(i,j)}(I(i,j)) = b_{(i,j)} + \frac{1}{2} I(i,j) \ln \frac{1-p}{p}$; we re-normalize all of the parameters such that $\max_{(i,j)} |\widetilde{b}_{(i,j)}(I(i,j))| = 1$.

We can modulate the difficulty of the inference problem with the noise level p as one would expect from the nature of the inference problem: because of the re-normalization, the magnitude w's of the resulting IM posterior increases with p. To see this, first note that the nature of our estimate of the prior biases is such that $|b_{(i,j)}| = 1$ because there is likely to be a pixel that is always black (-1) in the training images by the nature of the dataset. Also, the "observation" term in the expression of $\widetilde{b}_{(i,j)}$ decreases to 0 as $p \to 0.5$. Hence, at $p = 0.5$, the posterior IM is the prior. Certainly, the problem is becoming harder as p increases to 0.5. Yet, this increase in difficulty is not unfair to the standard methods because it is not harder than performing inference on the prior IM: if anything, the re-normalization makes inference in the posterior IM easier. We ran the exact same algorithms on the constructed Ising models as in the synthetic experiments, though we used a slightly different number of (max) iterations: 10^2 for fp (all variants); 10^4 for bp, nr, and mw (all variants); 10^5 for trw and gs; and 10^6 for mf.

Results. The top line of Fig. 2 shows that in fact the standard methods are unable to handle inference in the prior IM, as their estimate is that every pixel is black, while the CE-inspired methods do a more reasonable job: as expected from the nature of the images, the prior should reflect a signifi-

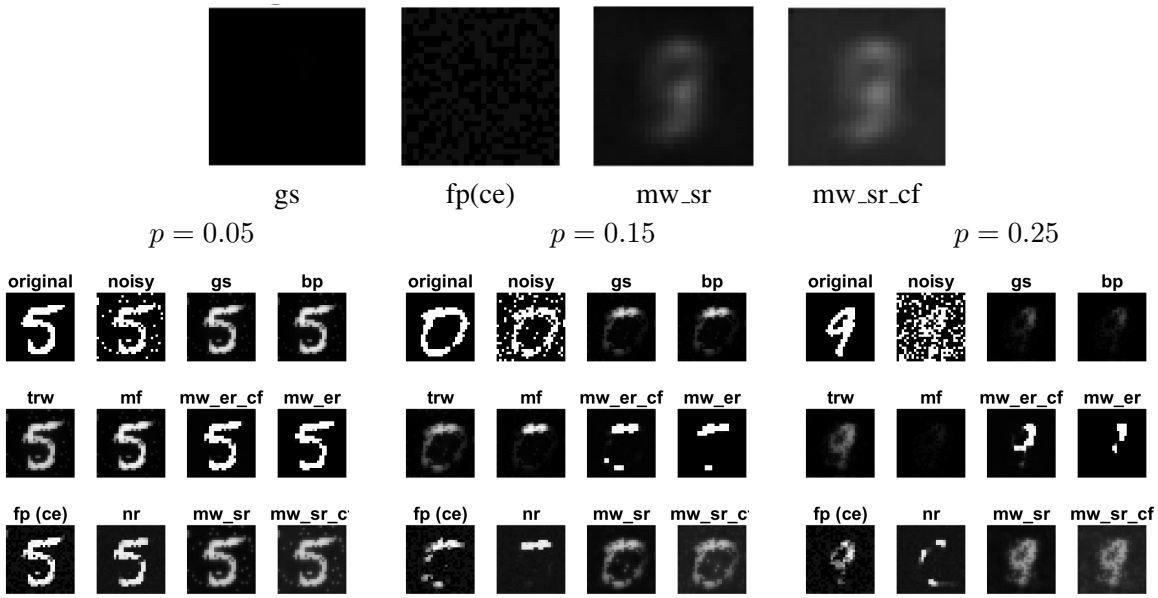

Figure 2: **Visualization of the Marginal Estimates for the MNIST-based IMs:** The top row shows grayscale images representing the estimates of the prior IM's marginal probability, for gs, fp(ce), mw_sr, and mw_sr_cf. The results for bp, mf, and trw are the same as that for gs in this case. The rest represents the results as the noise level p increases.

cantly more likely chance that the pixels near the center of the image are white. The "fuzziness" near the center of the image representing the marginal prior estimates for mw_sr and mw_sr_cf, which are CE-inspired methods, reflects that expectation. Those images also suggest that our albeit ad-hoc estimation of the IM parameters from the training images based on simple empirical averages for certain first- and second-moments is reasonable. (Note that MLE is intractable in this setting. The fuzzy shape looks like an "average" MNIST handwritten digit, a reasonable prior.) The bottom part of Fig. 2 presents a visualization of the algorithms' output at noise levels $p \in \{0.05, 0.15, 0.25\}$. The grayscale images show the corresponding marginal probability estimate for each individual pixel/node in the resulting posterior IM. Although we do include the original handwritten digit, it is not fair to compare the results directly to that original image: We are interested in probabilistic inference, while the original image is akin to a MAP assignment. The results for $p = 0.05$ are all reasonable, and somewhat consistent with the results of the synthetic experiments. As the noise level increases, the inference done by gs worsens, and the quality of bp, mf, nr, mw_er, and mw_er_cf deteriorates. At $p = 0.25$, the image becomes obscured to a point where even a human may not be able to confidently identify the original digit. Only trw, fp (ce), mw_sr, and mw_sr_cf maintain a reasonable quality of estimation at this noise level.

Closing Remark. We hope this work will start a conversation on the synergy between equilibrium computation and belief inference, which appears to be an intriguing and potentially fruitful research direction for both mathematical, algorithmic, and computational game-theory and PGMs.

Acknowledgements. This work was supported in part by NSF CAREER Award IIS-1643006 (transferred from IIS-1054541) and NSF Award IIS-1907553.

References

R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1, 1974.

A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In N. Nisan, T. Roughgarden, Éva Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 4, pages 79–102. Cambridge University Press, 2007.

J. Domke and X. Liu. Projecting Ising model parameters for fast mixing. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 665–673. 2013.

A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *STOC '04: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pages 604–612, New York, NY, USA, 2004. ACM.

D. Fudenberg and D. Levine. *The Theory of Learning in Games*. MIT Press, 1999.

J. Hammersley and P. Clifford. Markov fields on finite graphs and lattices. Unpublished, 1971.

F. Hamze and N. de Freitas. From fields to trees. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, UAI '04, pages 243–250, Arlington, Virginia, United States, 2004. AUAI Press.

J. Hannan. Approximation to Bayes risk in repeated play. In M. Dresher, A. W. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume III, pages 97–140. Princeton University Press, 1957.

S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127 – 1150, 2000.

J. Hofbauer and W. H. Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294, 2002.

A. X. Jiang and K. Leyton-Brown. Polynomial computation of exact correlated equilibrium in compact games. *SIGecom Exchanges*, 10(1):6–8, 2011.

A. X. Jiang and K. Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior*, 91:347 – 359, 2015.

S. Kakade, M. Kearns, J. Langford, and L. Ortiz. Correlated equilibria in graphical games. In *EC '03: Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 42–47, New York, NY, USA, 2003. ACM.

H. Kamisetty, E. P. Xing, and C. J. Langmead. Approximating correlated equilibria using relaxations on the marginal polytope. In L. Getoor and T. Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 1153–1160, New York, NY, USA, 2011. ACM.

Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.

R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1): 6–38, 1995.

D. Monderer and L. S. Shapley. Fictitious play property for games with identical interests. *Journal of Economic Theory*, 68(1):258 – 265, 1996.

H. Moulin and J. P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3):201–221, 1978.

J. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, September 1951.

J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America*, 35(1):48–49, Jan. 1950.

L. E. Ortiz. Graphical potential games. *CoRR*, abs/1505.01539, 2015. URL `http://arxiv.org/abs/1505.01539`.

L. E. Ortiz, R. E. Schapire, and S. M. Kakade. Maximum entropy correlated equilibria. In M. Meila and X. Shen, editors, *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics (AISTATS-07)*, volume 2, pages 347–354. Journal of Machine Learning Research - Proceedings Track, 2007.

C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In *STOC '05: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, pages 49–56, 2005.

C. H. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3):1–29, 2008.

I. Rezek, D. S. Leslie, S. Reece, S. J. Roberts, A. Rogers, R. K. Dash, and N. R. Jennings. On similarities between inference in game theory and machine learning. *J. Artif. Intell. Res. (JAIR)*, 33:259–283, 2008.

J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1947. Second Edition.

M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky. A new class of upper bounds on the log partition function. *IEEE Transactions on Information Theory*, 51(7):2313–2335, July 2005.