## Marginal polytopes of graphical models: Linear programs, max-product, and variational relaxation

Martin Wainwright<br>Department of Statistics<br>Department of Electrical Engineering and Computer Science<br>UC Berkeley, CA<br>Email: wainwrig@\{stat,eecs\}.berkeley.edu

Based on joint works with:

Tommi Jaakkola, Alan Willsky (MIT)<br>Michael Jordan (Univ. California, Berkeley)<br>Vladimir Kolmogorov (Univ. College London)<br>Alekh Agarwal, Pradeep Ravikumar (Univ. California, Berkeley)

## Introduction

- max/sum-product message-passing:
- "divide and conquer": based on factorization/Markov properties
- exact for decomposable; approximate for general graphs
- now standard in various fields (e.g., statistics, statistical machine learning, statistical physics, computer vision, computational biology....)
- convex relaxations (LP, SOCP, SDP etc.):
- "relax" a hard combinatorial problem into a simple convex one
- standard method in computer science, operations research, polyhedral combinatorics
- notion of marginal polytope:
- geometric object associated with any undirected graphical model
- complexity critically determined by graph topology
- yields fruitful connections between message-passing and LP relaxation


## MAP optimization in undirected graphical models



- undirected graph $G=(V, E)$
- $X_{s} \equiv$ random variable at node $s$ taking values $x_{s} \in \mathcal{X}_{s}$
- $\theta_{s}\left(x_{s}\right) \equiv$ observation term
- $\theta_{s t}\left(x_{s}, x_{t}\right) \equiv$ coupling term
- overall distribution decomposes additively on graph cliques:

$$
p(\mathbf{x} ; \theta) \propto \exp \left\{\sum_{s \in V} \theta_{s}\left(x_{s}\right)+\sum_{(s, t) \in E} \theta_{s t}\left(x_{s}, x_{t}\right)\right\}
$$

- mode or maximum a posteriori (MAP) estimate:

$$
\widehat{\mathbf{x}} \in \arg \max _{\mathbf{x} \in \mathcal{X}^{N}}\left\{\sum_{s \in V} \theta_{s}\left(x_{s}\right)+\sum_{(s, t) \in E} \theta_{s t}\left(x_{s}, x_{t}\right)\right\} .
$$

## Max-product on trees

Goal: Compute most probable configuration on a tree:

$$
\begin{gathered}
\widehat{\mathbf{x}}=\arg \max _{\mathbf{x} \in \mathcal{X}^{N}}\left\{\prod_{s \in V} \exp \left(\theta_{s}\left(x_{s}\right) \prod_{(s, t) \in E} \exp \left(\theta_{s t}\left(x_{s}, x_{t}\right)\right)\right\} .\right. \\
\max _{x_{1}, x_{2}, x_{3}} p(\mathbf{x})=\max _{x_{1}}\left[\exp \left(\theta_{1}\left(x_{1}\right)\right) \prod_{t \in\{1,3\}}\left\{\max _{x_{t}} \exp \left[\theta_{t}\left(x_{t}\right)+\theta_{2 t}\left(x_{2}, x_{t}\right)\right]\right\}\right]
\end{gathered}
$$

Max-product strategy: "Divide and conquer": break global maximization into simpler sub-problems. (Lauritzen \& Spiegelhalter, 1988; Dawid, 1992)

## Max-product recursions

Decompose: $\max _{x_{1}, x_{2}, x_{3}, x_{4}, x_{5}} p(\mathbf{x})=\max _{x_{1}}\left[\exp \left(\theta_{1}\left(x_{1}\right)\right) \prod_{t \in N(2)} M_{t 2}\left(x_{2}\right)\right]$.


Update messages:

$$
M_{32}\left(x_{3}, x_{2}\right)=\max _{x_{3}}\left[\exp \left(\theta_{3}\left(x_{3}\right)+\theta_{23}\left(x_{2}, x_{3}\right) \prod_{v \in N(3) \backslash 2} M_{v 3}\left(x_{3}\right)\right]\right.
$$

## Variational view: Max-product and linear programs

- MAP as integer program: $f^{*}=\max _{\mathbf{x} \in \mathcal{X}^{N}}\left\{\sum_{s \in V} \theta_{s}\left(x_{s}\right)+\sum_{(s, t) \in E} \theta_{s t}\left(x_{s}, x_{t}\right)\right\}$
- define local marginal distributions (e.g., for $m=3$ states):

$$
\mu_{s}\left(x_{s}\right)=\left[\begin{array}{l}
\mu_{s}(0) \\
\mu_{s}(1) \\
\mu_{s}(2)
\end{array}\right] \quad \mu_{s t}\left(x_{s}, x_{t}\right)=\left[\begin{array}{lll}
\mu_{s t}(0,0) & \mu_{s t}(0,1) & \mu_{s t}(0,2) \\
\mu_{s t}(1,0) & \mu_{s t}(1,1) & \mu_{s t}(1,2) \\
\mu_{s t}(2,0) & \mu_{s t}(2,1) & \mu_{s t}(2,2)
\end{array}\right]
$$

- alternative formulation of MAP as linear program

$$
g^{*}=\max _{\left(\mu_{s}, \mu_{s t}\right) \in \mathbb{M}(T)}\left\{\sum_{s \in V} \mathbb{E}_{\mu_{s}}\left[\theta_{s}\left(x_{s}\right)\right]+\sum_{(s, t) \in E} \mathbb{E}_{\mu_{s t}}\left[\theta_{s t}\left(x_{s}, x_{t}\right)\right]\right\}
$$

Local expectations: $\quad \mathbb{E}_{\mu_{s}}\left[\theta_{s}\left(x_{s}\right)\right]:=\sum_{x_{s}} \mu_{s}\left(x_{s}\right) \theta_{s}\left(x_{s}\right)$.

Key question: What constraints must local marginals $\left\{\mu_{s}, \mu_{s t}\right\}$ satisfy?

## Marginal polytopes for general undirected models

- $\mathbb{M}(G) \equiv$ set of all globally realizable marginals $\left\{\mu_{s}, \mu_{s t}\right\}$ :

$$
\left\{\vec{\mu} \in \mathbb{R}^{m^{N}} \mid \mu_{s}\left(x_{s}\right)=\sum_{x_{t}, t \neq s} p_{\mu}(\mathbf{x}), \text { and } \mu_{s t}\left(x_{s}, x_{t}\right)=\sum_{x_{u}, u \neq s, t} p_{\mu}(\mathbf{x})\right\}
$$

for some $p_{\mu}(\cdot)$ over $\left(X_{1}, \ldots, X_{N}\right) \in\{0,1, \ldots, m-1\}^{N}$.


- polytope in $m|V|+m^{2}|E|$ dimensions ( $m$ per vertex, $m^{2}$ per edge)
- with $m^{N}$ vertices
- number of facets?


## Marginal polytope for trees

- $\mathbb{M}(T) \equiv$ special case of marginal polytope for tree $T$
- local marginal distributions on nodes/edges (e.g., $m=3$ )

$$
\mu_{s}\left(x_{s}\right)=\left[\begin{array}{l}
\mu_{s}(0) \\
\mu_{s}(1) \\
\mu_{s}(2)
\end{array}\right] \quad \mu_{s t}\left(x_{s}, x_{t}\right)=\left[\begin{array}{lll}
\mu_{s t}(0,0) & \mu_{s t}(0,1) & \mu_{s t}(0,2) \\
\mu_{s t}(1,0) & \mu_{s t}(1,1) & \mu_{s t}(1,2) \\
\mu_{s t}(2,0) & \mu_{s t}(2,1) & \mu_{s t}(2,2)
\end{array}\right]
$$

Consequence of junction tree theorem: If $\left\{\mu_{s}, \mu_{s t}\right\}$ are nonnegative and locally consistent:

Normalization : $\quad \sum_{x_{s}} \mu_{s}\left(x_{s}\right)=1$
Marginalization: $\quad \sum_{x_{t}^{\prime}} \mu_{s t}\left(x_{s}, x_{t}^{\prime}\right)=\mu_{s}\left(x_{s}\right)$,
then on any tree-structured graph $T$, they are globally consistent. (Lauritzen \& Spiegelhalter, 1988)

## Max-product on trees: Linear program solver

- MAP problem as a simple linear program:

$$
f(\widehat{\mathbf{x}})=\arg \max _{\vec{\mu} \in \mathbb{M}(T)}\left\{\sum_{s \in V} \mathbb{E}_{\mu_{s}}\left[\theta_{s}\left(x_{s}\right)\right]+\sum_{(s, t) \in E} \mathbb{E}_{\mu_{s t}}\left[\theta_{s t}\left(x_{s}, x_{t}\right)\right]\right\}
$$

subject to $\vec{\mu}$ in tree marginal polytope:

$$
\mathbb{M}(T)=\left\{\vec{\mu} \geq 0, \quad \sum_{x_{s}} \mu_{s}\left(x_{s}\right)=1, \quad \sum_{x_{t}^{\prime}} \mu_{s t}\left(x_{s}, x_{t}^{\prime}\right)=\mu_{s}\left(x_{s}\right)\right\} .
$$

## Max-product and LP solving:

- on tree-structured graphs, max-product is a dual algorithm for solving the tree LP.
(Wai. \& Jordan, 2003)
- max-product message $M_{t s}\left(x_{s}\right) \equiv$ Lagrange multiplier for enforcing the constraint $\sum_{x_{t}^{\prime}} \mu_{s t}\left(x_{s}, x_{t}^{\prime}\right)=\mu_{s}\left(x_{s}\right)$.


## Standard message-passing algorithms: With cycles

Exact for trees, but approximate for graphs with cycles.


Question: What does max-product compute on a graph with cycles?

## Some previous theory on ordinary max-product

- optimal for trees, and junction trees (Lauritzen \& Spiegelhalter, 1988; Pearl, 1988; Dawid, 1992)
- analysis of graphs with large girth (Gallager, 1963; many others from 1990s onwards)
- single-cycle graphs (Aji \& McEliece, 1998; Horn, 1999; Weiss, 1998)
- existence of fixed points for positive couplings (Wainwright et al., 2003)
- local optimality guarantees:
-"tree-plus-loop" neighborhoods
(Weiss \& Freeman, 2001)
- strengthened optimality results and computable error bounds (Wainwright et al., 2003)
- some exactness results for particular types of matching problems (Bayati et al., 2006, 2008; Jebara \& Huang, 2007; Sanghavi, 2008)


## Standard analysis via computation tree

- standard tool: computation tree of message-passing updates (Gallager, 1963)

2

(a) Original graph

(b) Computation tree (4 iterations)

- level $t$ of tree: all nodes whose messages reach the root (node 1) after $t$ iterations of message-passing


## Illustration: Non-exactness of standard max-product

## Intuition:

- max-product solves (exactly) modified problem on computation tree
- edge/nodes not equally weighted $\Rightarrow$ incorrectness of max-product

(a) Diamond graph $G_{\text {dia }}$
- for example: asymptotic node fractions in this computation tree:

$$
\left[\begin{array}{llll}
f(1) & f(2) & f(3) & f(4)
\end{array}\right]=\left[\begin{array}{llll}
0.2393 & 0.2607 & 0.2607 & 0.2393
\end{array}\right]
$$

## A whole family of non-exact examples



$$
\begin{aligned}
& \theta_{s}\left(x_{s}\right)= \begin{cases}\alpha x_{s} & \text { if } s=1 \text { or } s=4 \\
\beta x_{s} & \text { if } s=2 \text { or } s=3\end{cases} \\
& \theta_{s t}\left(x_{s}, x_{t}\right)= \begin{cases}-\gamma & \text { if } x_{s} \neq x_{t} \\
0 & \text { otherwise }\end{cases}
\end{aligned}
$$

- for $\gamma$ sufficiently large, optimal solution is always either

$$
1^{4}=\left[\begin{array}{llll}
1 & 1 & 1 & 1
\end{array}\right] \text { or }(-1)^{4}=\left[\begin{array}{llll}
(-1) & (-1) & (-1) & (-1)
\end{array}\right]
$$

- max-product and optimal decision based on different boundaries:
$\underline{\text { Optimal boundary: }} \quad \widehat{\mathbf{x}}= \begin{cases}1^{4} & \text { if } 0.25 \alpha+0.25 \beta \geq 0 \\ (-1)^{4} & \text { otherwise }\end{cases}$
$\underline{\text { Max-product boundary: }} \quad \widehat{\mathbf{x}}= \begin{cases}1^{4} & \text { if } 0.2393 \alpha+0.2607 \beta \geq 0 \\ (-1)^{4} & \text { otherwise }\end{cases}$


## Tree-reweighted max-product algorithm

Message update from node $t$ to node $s$ :
$M_{t s}\left(x_{s}\right) \leftarrow \kappa \max _{x_{t}^{\prime} \in \mathcal{X}_{t}}\{\exp [\underbrace{\frac{\theta_{s t}\left(x_{s}, x_{t}^{\prime}\right)}{\rho_{s t}}}+\theta_{t}\left(x_{t}^{\prime}\right)] \frac{\prod_{v \in \mathcal{N}(t) \backslash s} \overbrace{\left[M_{v t}\left(x_{t}\right)\right]^{\rho_{v t}}}^{\left[M_{s t}\left(x_{t}\right)\right]^{\left(1-\rho_{t s}\right)}}}{\underbrace{[.} . . . . ~ . ~}$ reweighted edge opposite message

## Properties:

1. Modified updates remain distributed and purely local over the graph.

- Messages are reweighted with $\rho_{s t} \in[0,1]$.

2. Key differences:

- Potential on edge $(s, t)$ is rescaled by $\rho_{s t} \in[0,1]$.
- Update involves the reverse direction edge.

3. The choice $\rho_{s t}=1$ for all edges $(s, t)$ recovers standard update.

## Edge appearance probabilities

Experiment: What is the probability $\rho_{e}$ that a given edge $e \in E$ belongs to a tree $T$ drawn randomly under $\boldsymbol{\rho}$ ?

(a) Original

(b) $\rho\left(T^{1}\right)=\frac{1}{3}$

$$
\rho_{b}=1 ;
$$

(c) $\rho\left(T^{2}\right)=\frac{1}{3}$
(d) $\rho\left(T^{3}\right)=\frac{1}{3}$

In this example: $\quad \rho_{b}=1 ; \quad \rho_{e}=\frac{2}{3} ; \quad \rho_{f}=\frac{1}{3}$.

The vector $\boldsymbol{\rho}_{e}=\left\{\rho_{e} \mid e \in E\right\}$ must belong to the spanning tree polytope, denoted $\mathbb{T}(G)$.
(Edmonds, 1971)

## TRW max-product does not lie

- from message fixed point $M^{*}$, compute pseudo-max-marginals associated with vertex $s$,

$$
\nu_{s}\left(x_{s}\right)=\exp \left(\theta_{s}\left(x_{s}\right)\right) \prod_{t \in N(s)}\left[M_{t s}^{*}\left(x_{s}\right)\right]^{\rho_{t s}}
$$

and similar quantity for edge $(s, t)$.

- say strong tree agreement holds if there exists a configuration $\mathbf{x}^{*}$ such that:

$$
\begin{aligned}
x_{s}^{*} & \in \arg \max _{x_{s}} \nu_{s}\left(x_{s}\right) \quad \text { for all } s \in V \\
\left(x_{s}^{*}, x_{t}^{*}\right) & \in \arg \max _{x_{s}, x_{t}} \nu_{s t}\left(x_{s}, x_{t}\right) \quad \text { for all }(s, t) \in E .
\end{aligned}
$$

Theorem: For any fixed point $M^{*}$ any STA configuration $\mathbf{x}^{*}$ is a mode (most probable configuration) on the full graph $G$.
(WaiJaaWil05)

- sharp contrast to ordinary max-product, which does lie


## Tree-based relaxation for graphs with cycles

Set of locally consistent pseudomarginals for general graph $G$ :

$$
\mathbb{L}(G)=\left\{\vec{\tau} \mid \sum_{x_{s}} \tau_{s}\left(x_{s}\right)=1, \quad \sum_{x_{t}} \tau_{s t}\left(x_{s}, x_{t}^{\prime}\right)=\tau_{s}\left(x_{s}\right)\right\}
$$



Key: For a general graph, $\mathbb{L}(G)$ is an outer bound on $\mathbb{M}(G)$, and yields a linear-programming relaxation of the MAP problem:

$$
f(\widehat{\mathbf{x}})=\max _{\vec{\mu} \in \mathbb{M}(G)} \theta^{T} \vec{\mu} \leq \max _{\vec{\tau} \in \mathbb{L}(G)} \theta^{T} \vec{\tau}
$$

## TRW max-product and LP relaxation

First-order (tree-based) LP relaxation:

$$
f(\widehat{\mathbf{x}}) \leq \max _{\vec{\tau} \in \mathbb{L}(G)}\left\{\sum_{s \in V} \mathbb{E}_{\tau_{s}}\left[\theta_{s}\left(x_{s}\right)\right]+\sum_{(s, t) \in E} \mathbb{E}_{\tau_{s t}}\left[\theta_{s t}\left(x_{s}, x_{t}\right)\right]\right\}
$$

Theorem: (WaiJaaWil05; Kolmogorov \& Wainwright, 2005):
(a) Strong tree agreement Any TRW fixed-point that satisfies the strong tree agreement condition specifies an optimal LP solution.
(b) LP solving: For any binary pairwise problem, TRW max-product solves the first-order LP relaxation.
(c) Persistence for binary problems: Let $S \subseteq V$ be the subset of vertices for which there exists a single point $x_{s}^{*} \in \arg \max _{x_{s}} \nu_{s}^{*}\left(x_{s}\right)$. Then for any optimal solution, it holds that $y_{s}=x_{s}^{*}$.

## Basic idea: convex combinations of trees

Observation: Easy to find its MAP-optimal configurations on trees:

$$
\operatorname{OPT}(\theta(T)):=\left\{\mathbf{x} \in \mathcal{X}^{n} \mid \mathbf{x} \text { is MAP-optimal for } p(\mathbf{x} ; \theta(T))\right\}
$$

Idea: Approximate original problem by a convex combination of trees.

$$
\begin{array}{ll}
\rho=\{\rho(T)\} & \equiv \text { probability distribution over spanning trees } \\
\theta(T) & \equiv \text { tree-structured parameter vector }
\end{array}
$$



## Dual perspective: linear programming relaxation

- Upper bound maintained by reweighted message-passing:

$$
\max _{\mathbf{x} \in \mathcal{X}^{N}}\left\langle\theta^{*}, \phi(\mathbf{x})\right\rangle \leq \sum_{T \in \mathfrak{T}} \rho(T) \max _{\mathbf{x} \in \mathcal{X}^{N}}\langle\theta(T), \phi(\mathbf{x})\rangle
$$

- Dual of finding optimal upper bound $\equiv$ tree-based LP relaxation:

$$
\max _{\mathbf{x} \in \mathcal{X}^{N}}\left\langle\theta^{*}, \phi(\mathbf{x})\right\rangle \leq \max _{\mu \in \operatorname{LOCAL}(G)}\langle\mu, \phi(\mathbf{x})\rangle
$$

- TRW-MP algorithm fixed points specify LP optimum:
- whenever strong tree agreement holds
(WaiJaaWil05)
- for any binary problem
(KolWai05)
- ....but TRW-MP does not solve LP in general


## Various connections and extensions

- max-sum diffusion framework (Schlesinger et al., 1960s, 70s; Werner, 2007)
- binary QPs and roof duality: equivalent to relaxation using $\mathbb{L}(G)$ (Hammer et al., 1984; Boros et al., 1990)
- hierarchy of LP relaxations based on treewidth:

$$
\mathbb{M}(G)=\mathbb{L}_{t}(G) \subset \mathbb{L}_{t-1}(G) \subset \ldots \subset \mathbb{L}_{1}(G)
$$

- treewidth hierarchy: equivalent to Boros et al. (1990) and Sherali-Adams (1990) hierarchies for binary problems
- other approaches with links to first-order $\mathbb{L}(G)$ LP relaxation:
- sequential TRW and conv. guarantees
(Kolmogorov, 2005)
- convex free energies
(Weiss et al., 2007)
- sub-gradients
(Feldman et al, 2003; Komodakis et al., 2007)
- proximal projections
(Ravikumar et al., 2008)


## Extensions to computing/bounding likelihoods

- $\log$ normalization/likelihood for an undirected model:

$$
A(\theta)=\log \sum_{\mathbf{x} \in \mathcal{X}^{N}} \exp \left\{\sum_{s \in V} \theta_{s}\left(x_{s}\right)+\sum_{(s, t) \in E} \theta_{s t}\left(x_{s}, x_{t}\right)\right\}
$$

- variational reformulation as a convex optimization problem:

$$
A(\theta)=\max _{\vec{\mu} \in \mathbb{M}(G)}\left\{\theta^{T} \vec{\mu}+H(\vec{\mu})\right\}
$$

where

- $H(\vec{\mu})$ is maximized entropy, over all distributions with mean parameters $\vec{\mu}$
- marginal polytope $\mathbb{M}(G)$ of all globally realizable distributions
- both $H(\cdot)$ and $\mathbb{M}(G)$ pose significant challenges for general graphs
- as before hypertrees are easy, and inspire the same relaxation philosophy


## Summary

- marginal polytope: fundamental object associated with any discrete graphical model
- connections between LP relaxation and message-passing algorithms on graphs
- marginal polytopes and relaxations: also relevant for approximating/bounding marginals and likelihoods
- many open questions/issues:
- approximation guarantees for LP relaxations: role of graph structure
- guarantees for marginal/likelihood approximations
- extensions to mixed discrete/continuous graphs, non-parametric settings
- hybrid variational and MCMC methods


## Some papers

- Wainwright, M. J. \& Jordan, M. (2003) Graphical models, exponential families, and variational methods. Department of Statistics, UC Berkeley, Technical Report 649. To appear in Foundation and Trends in Machine Learning.
- Wainwright, M. J., Jaakkola, T. S., and Willsky, A. S., (2005), Exact MAP estimates via agreement on hypertrees: Message-passing and linear programming. IEEE Trans. Information Theory, 51:3697-3717.
- Wainwright, M. J., Jaakkola, T. S. and Willsky, A. S. (2005). A new class of upper bounds on the log partition function. IEEE Transactions on Information Theory. July, 51:2313-2335.
- Daskalakis, C., Dimakis, A. D., Karp, R. and Wainwright, M. J. (2008). Probabilistic analysis of linear programming decoding. To appear in IEEE Trans. Info. Theory.
- Ravikumar, P., Agarwal, A. and Wainwright, M. J. (2008). Message-passing for graph-structured linear programs: Proximal projections and convergence. To appear in Int. Conference on Machine Learning, Helsinki, Finland.

