
Randomized Optimum Models for Structured Prediction

Daniel Tarlow
University of Toronto
Dept. of Computer Science

Ryan P. Adams
School of Engineering & Applied Sciences
Harvard University

Richard S. Zemel
University of Toronto
Dept. of Computer Science

Abstract

One approach to modeling structured discrete data is to describe the probability of states via an energy function and Gibbs distribution. A recurring difficulty in these models is the computation of the partition function, which may require an intractable sum. However, in many such models, the mode can be found efficiently even when the partition function is unavailable. Recent work on Perturb-and-MAP (PM) models (Papandreou and Yuille, 2011) has exploited this discrepancy to approximate the Gibbs distribution for Markov random fields (MRFs). Here, we explore a broader class of models, called Randomized Optimum models (RandOMs), which include PM as a special case. This new class of models encompasses not only MRFs, but also other models that have intractable partition functions yet permit efficient mode-finding, such as those based on bipartite matchings, shortest paths, or connected components in a graph. We develop likelihood-based learning algorithms for RandOMs, which, empirical results indicate, can produce better models than PM.

1 Introduction

One of the fundamental frustrations when modeling high-dimensional phenomena is that it is often easy to *score* configurations of data (e.g., using an energy function), but normalizing the distribution that this score implies may require enumerating a very large space (e.g., computing a partition function). When the normalizing constant is unknown, it becomes difficult to perform learning and inference, as one cannot assess which scoring functions give higher probability to the data. This can be somewhat alleviated when exact

data can be generated from the model (Murray et al., 2006; Adams et al., 2009), but such situations are rare.

A curious computational phenomenon is that there are important and useful classes of energy-based models in which summing over all states is difficult, but *optimizing* is easy. That is, there are models in which the Gibbs-based probabilistic interpretation yields an unnormalizable model, but the configuration that minimizes the energy function (the mode, or maximum a posteriori (MAP) configuration) can be found exactly and efficiently. Two common examples of this are the graph cut algorithm for finding the mode of an associative (submodular) binary Markov random field (Kolmogorov and Zabih, 2004), and the bipartite matching problem, in which solutions can be found using the Hungarian algorithm (Munkres, 1957) in $\mathcal{O}(n^3)$ time, but computing the normalizing constant is $\#P$ -complete (Valiant, 1979).

The complexity contrast between optimizing and normalizing has led to two different styles of learning in models of high-dimensional data. The first relies on optimization, but dispenses with a probabilistic interpretation, instead seeking to ensure that the ground truth has the optimal score. The energy-based models championed by LeCun et al. (2006), and structured SVMs trained via maximum margin methods (Tsochantaridis et al., 2005) both fall into this category. The second approach retains the probabilistic formulation, but utilizes approximate inference, such as loopy belief propagation, pseudo-likelihood, or sampling, to approximate the enumeration that is required for computing the probability of the data.

Our aim is to bridge this gap, developing probabilistic models that avoid sampling over complex discrete spaces, and the approximations required by enumeration. Recent work by Papandreou and Yuille (2011) has taken a step in this direction, in the particular case of Markov random fields. These authors reformulate a Gibbsian probabilistic model as injecting local noise into the energy function followed by mode-finding under the perturbed objective. They call this approach *Perturb-and-MAP*, and consider factorized noise injection

Appearing in Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS) 2012, La Palma, Canary Islands. Volume 22 of JMLR: W&CP 22. Copyright 2012 by the authors.

tion based on the Gumbel distribution, which aims to approximate the Gibbs distribution implied by the original Markov random field.

In this paper, we develop a broader class of *Randomized Optimum models* (RandOMs), which are probabilistic models that exploit efficient optimization. RandOMs generalize previous work along three different axes. First, this new approach can be used to develop probabilistic models for complex structured spaces that are not typically modeled probabilistically, but in which global optima can be computed efficiently. We prove representational limitations to standard random field models and show that RandOMs can be used in a case where the standard graphical modeling representation fails. Second, we disconnect from attempting to approximate the Gibbs distribution and explore other ways of moving from an energy function to a probability distribution over structured outputs. Third, we develop new maximum likelihood-based learning schemes with different tradeoffs between efficiency and representation of uncertainty.

In total, we continue further in the direction of developing models that allow a better balance of probabilistic rigor with computational tractability. We see three main benefits to the approach: (a) it allows proper probabilistic modeling of structured domains: support is only given to configurations that satisfy domain-specific constraints; (b) an exact sample can be drawn from the models in polynomial time, without the need for a Markov chain; and (c) we believe our novel learning approaches to be a promising direction forward for learning with discrete output spaces that are more complex than are typically addressed by structured output learning. A key point related to the tractability of our approach is that we will map the complex discrete space into a well-behaved continuous latent space, then perform the difficult computations in the continuous space, where we can leverage powerful computational tools and concepts.

2 Randomized Optimum models

The goal of the Randomized Optimum model is to learn families of conditional distributions over highly structured spaces. The centerpiece of this approach is to specify these distributions in terms of the minima of randomized objective functions.

Let \mathcal{Y} define a structured space of interest, such as the space of bipartite matchings, paths on a graph, or segmentations of a graph. An element $\mathbf{y} \in \mathcal{Y}$ is represented as the joint assignment of D categorical variables, $\mathbf{y} = \{y_1, \dots, y_D\}$, where $y_d \in L_d = \{1, \dots, K_d\}$, and K_d is the cardinality of variable d . In the supervised setting that we explore here, we wish to construct a distribution on this space by conditioning upon a




Figure 1: Overview of model. Starting from right to left: an observation of structured output \mathbf{y} determines the inverse mapping set $F^{-1}(\mathbf{y})$, which is denoted by the shaded region. We then have a region in \mathbf{w} space that leads to \mathbf{y} after a discrete optimization algorithm is run. ψ and $\phi(\mathbf{x})$ define a distribution over \mathbf{y} via $p(\mathbf{w} | \psi, \phi(\mathbf{x}))$.

feature space \mathcal{X} . The approach we take is to mediate the interaction between \mathcal{X} and \mathcal{Y} via real-valued latent variables $\mathbf{w} \in \mathcal{W} = \mathbb{R}^P$, which we will discuss in more detail below. As an example, for paths on a graph, features may be properties of the edges and vertices.

An overview of the model is shown in Fig. 1. There are two main components of the model: (a) the mapping from features to a distribution in latent variable space, $p(\mathbf{w} | \psi, \phi(\mathbf{x}))$; (b) the mapping from latent variable space to observations \mathbf{y} . In a standard CRF formulation, the (a) is deterministic while (b) is random. In the RandOM formulation, these are reversed: (a) is random, while (b) is deterministic. We discuss the relationship to standard CRF models more thoroughly in the Appendix, showing specifically how to map standard CRF notation into our formulation.

We will assume that there exists some matrix function $\phi : \mathcal{X} \rightarrow \mathbb{R}^{F \times P}$ such that model parameters $\psi \in \mathbb{R}^F$ can be used in conjunction with $\phi(\mathbf{x})$ to define a distribution over \mathbf{w} , which will then be used to define a distribution over \mathbf{y} . Our objective is to find parameters ψ that most effectively specify a family of distributions $p(\mathbf{y} | \psi, \phi(\mathbf{x})) = \int_{\mathbf{w}} p(\mathbf{y} | \mathbf{w}) p(\mathbf{w} | \psi, \phi(\mathbf{x})) d\mathbf{w}$, where $\mathbf{y} \in \mathcal{Y}$, given labeled data $\{\mathbf{x}_n, \mathbf{y}_n\}_{n=1}^N$.

In Randomized Optimum models, we define these distributions via a family of objective functions $f_{\mathbf{w}} : \mathcal{Y} \rightarrow \mathbb{R}$ which are indexed by the latent variables \mathbf{w} . Then, \mathbf{y} comes not from a distribution given \mathbf{w} but is the deterministic optimum of $f_{\mathbf{w}}(\mathbf{y})$. We assume that there is a unique minimum for $f_{\mathbf{w}}(\mathbf{y})$ for each \mathbf{w} in \mathcal{W} , yielding a simple generative model:

$$\mathbf{w} \sim p(\mathbf{w} | \psi, \phi(\mathbf{x})) \quad \mathbf{y} = \underset{\mathbf{y}'}{\operatorname{argmin}} f_{\mathbf{w}}(\mathbf{y}').$$

Given the features \mathbf{x} and the model parameters ψ , a random objective function is chosen via \mathbf{w} . The observed data are then specified as the unique minimum of this objective. We will denote this minimization as an implicit function $F(\mathbf{w}) = \underset{\mathbf{y}'}{\operatorname{argmin}} f_{\mathbf{w}}(\mathbf{y}')$, which has domain \mathcal{W} and codomain \mathcal{Y} . We also make heavy

use of the inverse set $F^{-1}(\mathbf{y})$, where:

$$\mathbf{w} \in F^{-1}(\mathbf{y}) \Rightarrow F(\mathbf{w}) = \mathbf{y} \quad (1)$$

The core *probabilistic* component of this model is the conditional distribution $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$. In the generative model, this is the distribution that determines which optimization problem will be solved to find the dependent variable \mathbf{y} . Many regression models could be used for RandOMs, e.g., Hannah et al. (2011); however, for concreteness, the reader may think of the model in terms of simple basis function regression with spherical Gaussian noise and the mean constrained to the unit hypercube by the logistic function:

$$\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}), \nu \sim \mathcal{N}(\sigma(\boldsymbol{\psi}^\top \phi(\mathbf{x})), \nu \mathbb{I}) \quad (2)$$

where $\sigma(u) = (1 + \exp\{-u\})^{-1}$. We will always further assume that \mathbf{w} is restricted to the unit hypercube.

2.1 Optimization Component of Model

Having specified the rest of the general framework, we turn our attention to the specification of $f_{\mathbf{w}}(\mathbf{y})$. In this paper we will focus on a set of specific cases in which Randomized Optimum models offer interesting computational advantages over alternative ways to specify families of distributions over structured spaces. There are three main requirements when choosing the objective function $f_{\mathbf{w}}(\mathbf{y})$. First, $f_{\mathbf{w}}(\mathbf{y})$ should enforce global constraints. For example, in the bipartite matching domain, $f_{\mathbf{w}}$ should enforce the constraint that each point is matched to exactly one other point. Second, the parameterization \mathbf{w} should give enough flexibility to model the domain of interest. By modifying \mathbf{w} , we should be able to prefer favorable structured objects over unfavorable structured objects. Third, the function $\arg \min_{\mathbf{y}} f_{\mathbf{w}}(\mathbf{y})$ must be efficient to compute.

There are many possible functions $f_{\mathbf{w}}(\mathbf{y})$ that satisfy these requirements for various structured \mathcal{Y} . We consider three classes of such functions, noting that not all function within these classes will be suitable:

A. Standard exponential family models, where

$$f_{\mathbf{w}}(\mathbf{y}) = \langle \mathbf{w}, \rho(\mathbf{y}) \rangle, \quad (3)$$

where $\rho(\mathbf{y})$ is the standard sufficient statistic in the exponential family representation of a MRF. This is the example considered by Papandreou and Yuille (2011), specifically in the case of a ferromagnetic Ising model. This standard CRF model supports efficient optimization via graph cuts.

B. Exponential family models with combinatorial base measure:

$$f_{\mathbf{w}}(\mathbf{y}) = \langle \mathbf{w}, \rho(\mathbf{y}) \rangle + \eta(\mathbf{y}). \quad (4)$$

The function $\eta(\mathbf{y})$ specifies the support constraints for the model. That is, it assigns infinite cost in the objective function to values of \mathbf{y} outside the desired sample space. For example $\eta(\mathbf{y})$ might encode the requirement that valid optima for $f_{\mathbf{w}}(\mathbf{y})$ be one-to-one matches in the bipartite matching problem. In the exponential family view, this can be thought of as the negative log of the base measure, which would take value infinity for values of \mathbf{y} which should have zero probability. Note that representing these constraints within a graphical model formulation often require high order potentials and can be difficult to work with.

C. Beyond exponential families. These are functions $f_{\mathbf{w}}(\mathbf{y})$ that provably cannot be expressed in a standard graphical model formulation, even when using high order potentials, yet still admit efficient optimization. We give an example, make the above claim precise, and provide a proof in the Appendix.

This formulation allows a wide range of Randomized Optimum models with useful properties on structured spaces. In the remainder of this section, we examine two such cases of interest. Two other examples are given in the Appendix.

2.1.1 Example: Bipartite Matching

In the bipartite matching problem, we have an undirected bipartite graph $G = (\mathcal{V}, \mathcal{E})$, with equal-sized partite sets $\mathbf{A} \subseteq \mathcal{V}$ and $\mathbf{B} \subseteq \mathcal{V}$ i.e., $\mathbf{A} \cap \mathbf{B} = \emptyset$, $\mathbf{A} \cup \mathbf{B} = \mathcal{V}$, and $J = |\mathbf{A}| = |\mathbf{B}|$. Associated with each pair (or edge) $(a_i \in \mathbf{A}, b_j \in \mathbf{B})$ is a real value w_{ij} , which gives a cost incurred for matching a_i to b_j . The objective of bipartite matching is to choose a minimum cost set of pairs $\{(a_{i_1}, b_{j_1}), \dots, (a_{i_{|\mathbf{A}|}}, b_{j_{|\mathbf{B}|}})\}$ such that each a_i and each b_j is paired with exactly one element of the opposing partite set. When (a_i, b_j) is included in the matching, we say that a_i is *matched* to b_j .

Bipartite matching enforces the concept of 1-to-1 correspondence, so it is not surprising that this structure arises often in real world problems in a variety of domains. Some notable examples include the problem of modeling correspondences in computer vision (Torresani et al., 2008); the assignment of workers to jobs in operations research; or computing word alignment in natural language processing (Taskar et al., 2005).

To express the bipartite matching problem in RandOM terms, we let $\mathbf{y} = \{y_{11}, \dots, y_{JJ}\}$ be a length- J^2 boolean vector, where y_{ij} denotes the event that a_i is matched to b_j . In this case, the sufficient statistic is simply the identity: $\rho(\mathbf{y}) = \mathbf{y}$. Then $\mathbf{w} = \{w_{11}, \dots, w_{JJ}\}$ is a length- J^2 real-valued vector, where w_{ij} gives the cost for matching a_i to b_j . Finally, $\eta(\mathbf{y})$ enforces the one-to-one matching constraint by assigning infinite cost to any \mathbf{y} that is not a valid matching. It should then be clear that

$$f_{\mathbf{w}}(\mathbf{y}) = \langle \mathbf{w}, \rho(\mathbf{y}) \rangle + \eta(\mathbf{y}) = \sum_{ij} \mathbf{y}_{ij} \mathbf{w}_{ij} + \eta(\mathbf{y}), \quad (5)$$

and so minimizing $f_{\mathbf{w}}$ is equivalent to the minimum cost bipartite matching problem.

2.1.2 Example: Connected Components

There are many scenarios where observations may come in the form of a graph-structured clustering of points. For example, image segmentation can be thought of as a clustering of pixels over a lattice graph. One possible model of these outputs is as connected components, as in (Turaga et al., 2010). A natural goal is to define a proper probabilistic model of outputs generated by such a connected components procedure. This leads to an example of the third general type — that which cannot be expressed as a standard graphical model, even with high order terms. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with non-negative edge weights $\mathbf{w} = \{w_{ij} : (i, j) \in \mathcal{E}\}$ and a fixed threshold τ , we wish the minimum cost solution to be to label variables according to the connected component that they belong to, where an edge is deemed to exist if and only if its corresponding weight is above the threshold i.e. $w_{ij} > \tau$. To represent this labeling, we define variables $\mathbf{y} = \{y_1, \dots, y_{|\mathcal{V}|}\} \in \{1, \dots, |\mathcal{V}|\}^{|\mathcal{V}|}$, and force the label of all variables within a component to be the minimum index of any variable within that component. Under this definition, conditional on a given \mathcal{G} and \mathbf{w} , there is a single desired labeling of variables, and all other labelings are equally undesirable.

More formally, we can describe $f_{\mathbf{w}}(\mathbf{y})$ using the following definitions:

$$\mathcal{P}(i, j) \text{ is the set of all paths from } i \text{ to } j \text{ in } \mathcal{G} \quad (6)$$

$$b_{ij}(\mathbf{w}) = \max_{p \in \mathcal{P}(i, j)} \min_{(l, m) \in p} w_{lm} \quad (7)$$

$$\mathcal{C}_i(\mathbf{w}) = \{j : i \neq j, b_{ij}(\mathbf{w}) > \tau\} \quad (8)$$

$$f_{\mathbf{w}}(\mathbf{y}) = \sum_{(i, j) \in \mathcal{E}} \mathbf{1}_{\{y_i = y_j\}} \max(0, b_{ij}(\mathbf{w}) - \tau) \quad (9)$$

$$+ \sum_{(i, j) \in \mathcal{E}} \mathbf{1}_{\{y_i \neq y_j\}} \max(0, \tau - b_{ij}(\mathbf{w})) + \eta(\mathbf{y})$$

$$\eta(\mathbf{y}) = \lim_{a \rightarrow \infty} a \cdot \left(1 - \prod_{i \in \mathcal{V}} [y_i = \min_{j \in \mathcal{C}_i(\mathbf{w})} j] \right). \quad (10)$$

It should be clear that by assigning \mathbf{y} as the result of a connected components algorithm, $f_{\mathbf{w}}(\mathbf{y}) = 0$. All other assignments have $f_{\mathbf{w}}(\mathbf{y}) > 0$, so we can exactly and efficiently compute the arg-minimum of this $f_{\mathbf{w}}$. An interesting question is whether there exists an exponential family representation (i.e., in the form of Eq. 4) of this model. We will prove that this is impossible via the geometry of the inverse set $F^{-1}(\mathbf{y})$.

3 Inverse Set Geometry

We believe a productive means of thinking about RandOM models is via the geometry of the inverse mapping set $F^{-1}(\mathbf{y})$. It is this geometry that determines the tractability of a RandOM as $f_{\mathbf{w}}$ becomes more complex. Due to space constraints, discussion of important geometric properties, along with proofs, will be discussed in the Appendix. Here, we simply state the important results:

Proposition 1. *When $f_{\mathbf{w}}(\mathbf{y})$ is defined as the exponential family (possibly with combinatorial base measure) as in Eq. 4, $F^{-1}(\mathbf{y})$ is a convex set.*

Proposition 2. *The inverse set $F^{-1}(\mathbf{y})$ for the connected components problem from Section 2.1.2 is a star-convex set.*

Lemma 1. *Let $f_{\mathbf{w}}(\mathbf{y})$ be defined as the connected components objective in Eq. 9. There is no equivalent expression of $f_{\mathbf{w}}(\mathbf{y})$ in the exponential family form Eq. 4.*

4 Learning RandOMs

The maximum likelihood learning objective is to maximize the amount of probability mass assigned to the ground truth labeling by the model. In the RandOM formulation, because the relationship between latent variables \mathbf{w} and observations \mathbf{y} is deterministic, this amounts to placing as much mass of $p(\mathbf{w} | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}))$ as possible inside the $F^{-1}(\mathbf{y})$ sets. In this section, we design learning algorithms for RandOMs that take advantage of the structure of $F^{-1}(\mathbf{y})$.

The principal challenge we need to address is how to handle the unknown latent variable \mathbf{w} for each training case. For any given observed configuration \mathbf{y} , there is a set of \mathbf{w} corresponding to objective functions for which $F(\mathbf{w}) = \mathbf{y}$. Recall that this set is the inverse mapping set $F^{-1}(\mathbf{y})$. The likelihood of $\boldsymbol{\psi}$ requires integrating over this set:

$$p(\{\mathbf{y}_n\}_{n=1}^N | \boldsymbol{\psi}, \{\mathbf{x}_n\}_{n=1}^N) = \prod_{n=1}^N \int_{F^{-1}(\mathbf{y}_n)} p(\mathbf{w}_n | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}_n)) d\mathbf{w}_n.$$

Because the optimization procedure is a many-to-one mapping, and since exactly integrating over $F^{-1}(\mathbf{y})$ is intractable (even in the case where $F^{-1}(\mathbf{y})$ is a convex set and $p(\mathbf{w}_n | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}_n))$ is Gaussian), there are essentially two practical approaches: approximately integrate over \mathbf{w} using Markov chain Monte Carlo (MCMC); or find a point estimate for \mathbf{w} by solving a constrained optimization problem. In this section, we will discuss formulations of both of these approaches where efficient combinatorial algorithms can be used to make learning more efficient. An orthogonal issue that we will not discuss further is the treatment of $\boldsymbol{\psi}$. Here, we will find a point estimate, but it would be straightforward to integrate over it via sampling.

4.1 Monte Carlo Expectation Maximization

Although we cannot generally analytically integrate over the pre-image set, we can nevertheless instantiate the latent variables $\{\mathbf{w}_n\}_{n=1}^N$ explicitly in a Markov chain, drawing samples using a specialized MCMC procedure. This general approach has the property that it allows for maximal flexibility in the choice of $p(\mathbf{w} | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}))$, enabling any option that would be available in the case where the \mathbf{w} themselves were the observations, such as nonparametric mixture models. We emphasize that sampling over $F^{-1}(\mathbf{y})$ is very different from the typical sampling over a complex discrete space.

More specifically, we write the joint distribution $p(\boldsymbol{\psi}, \{\mathbf{w}_n, \mathbf{y}_n, \mathbf{x}_n\}_{n=1}^N)$ as,

$$p(\boldsymbol{\psi}) \prod_{n=1}^N p(\mathbf{w}_n | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}_n)) \mathbf{1}_{\{\mathbf{w}_n \in F^{-1}(\mathbf{y}_n)\}} \quad (11)$$

Note that $p(\boldsymbol{\psi}, \{\mathbf{w}_n, \mathbf{y}_n, \mathbf{x}_n\}_{n=1}^N)$ has two parts. The first part, $p(\boldsymbol{\psi}) \prod_{n=1}^N p(\mathbf{w}_n | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}_n))$, is any standard regression model for \mathbf{w}_n given inputs \mathbf{x}_n , using parameters $\boldsymbol{\psi}$. The non-standard part of this formulation comes from the second part — the indicator function $\mathbf{1}_{\{\mathbf{w}_n \in F^{-1}(\mathbf{y}_n)\}}$, which constrains each \mathbf{w}_n to lie in the inverse mapping set $F^{-1}(\mathbf{y}_n)$ by assigning zero probability to any \mathbf{w}_n that lies outside the set. One of the main operations that we will need to repeatedly perform during learning is to check whether a given \mathbf{w}_n lies within $F^{-1}(\mathbf{y}_n)$. This “set membership” query could be done simply by running the optimization algorithm associated with the RandOM. We will discuss a much more efficient strategy, and its tailoring specifically to slice sampling, in Section 4.1.1.

Monte Carlo expectation maximization (MCEM) (Wei and Tanner, 1990) alternates between computing a Monte Carlo estimate of the expected complete data log likelihood (the E-step) and maximizing $\boldsymbol{\psi}$ given this approximation (the M-step). In MCEM iteration t , we use the procedure described in Section 4.1.1 to generate K samples of the latent variables \mathbf{w}_n using the current parameter estimate $\boldsymbol{\psi}^{(t)}$: $\mathbf{w}_{n,k}^{(t)} \sim p(\mathbf{w}_n | \boldsymbol{\psi}^{(t)}, \mathbf{x}_n) \cdot \mathbf{1}_{\{\mathbf{w}_n \in F^{-1}(\mathbf{y}_n)\}}$. These samples are then used to estimate the complete data log likelihood:

$$\begin{aligned} \mathbb{E}_{\mathbf{w}^{(t)}} [\ln p(\{\mathbf{y}_n, \mathbf{w}_n\}_{n=1}^N | \boldsymbol{\psi}, \{\mathbf{x}_n\}_{n=1}^N)] \\ \approx \sum_{n=1}^N \ln \left\{ \frac{1}{K} \sum_{k=1}^K p(\mathbf{w}_{n,k}^{(t)} | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}_n)) \right\}. \end{aligned} \quad (12)$$

This Monte Carlo approximation to the E-step is then used in the M-step to find the new parameter settings:

$$\boldsymbol{\psi}^{(t+1)} = \operatorname{argmax}_{\boldsymbol{\psi}} \sum_{n=1}^N \ln \left\{ \frac{1}{K} \sum_{k=1}^K p(\mathbf{w}_{n,k}^{(t)} | \boldsymbol{\psi}, \boldsymbol{\phi}(\mathbf{x}_n)) \right\}.$$

In each E-step, the Markov chain for each latent variable is initialized to its ending value from the previous step. As the inverse mapping does not change during updates to $\boldsymbol{\psi}$, the new starting location for \mathbf{w}_n is guaranteed to lie within $F^{-1}(\mathbf{y}_n)$. In the experimental section, we explore the effect of varying the number of samples to draw in each E-step, considering the case of using as few as 1 sample per instance per iteration.

4.1.1 Specialized Slice Sampler

Critical to MCEM (and fully-Bayesian inference) is the task of sampling the latent \mathbf{w} . In this section, we will discuss why this problem is particularly well suited to slice sampling (Neal, 2003), and we will show how dynamic combinatorial algorithms and the inner workings of slice sampling can be made to work synergistically together. For a detailed review of slice sampling, we recommend Neal (2003). We assume familiarity of basic slice sampling, and give a detailed description of our implementation of slice sampling, specialized for the case of bipartite matching models, in Algorithm 2 (see Appendix). The same principles apply for other RandOMs. Here we discuss the non-standard properties of our specialized slice sampler.

First, we maintain algorithmic state $\mathcal{A}_{\mathbf{y}}$ throughout the execution of the sampling algorithm. Whenever we make a set membership query (IN-INVERSE-SET in Alg. 2), we warm-start from the solution of the previous call, using dynamic combinatorial optimization e.g. for bipartite matching problems, we use the dynamic Hungarian algorithm of Korsah et al. (2007).

Second, we re-sample blocks of \mathbf{w} so as to be compatible with the dynamic combinatorial algorithm. In the bipartite matching case, we resample a subset of dimensions of \mathbf{w} at a time, where the subset corresponds to the terms that define the costs for a single row of the bipartite matching cost matrix. With this choice, the dynamic Hungarian algorithm can check inverse set membership in worst case $\mathcal{O}(J^2)$ time rather than the worst case $\mathcal{O}(J^3)$ time that would be required if all entries of \mathbf{w} changed. More importantly, empirically, the membership checks after a single row update often behave as $\mathcal{O}(1)$ or $\mathcal{O}(J)$, depending on how many values of $F(\mathbf{w})$ change as a result of the modification. If we modify \mathbf{w} without leaving the $F^{-1}(\mathbf{y})$ set, then by definition, $F(\mathbf{w})$ has not changed, making the dynamic algorithm very fast. When the sampler does step outside $F^{-1}(\mathbf{y})$, it is typically by a small amount, leading to a new $F(\mathbf{w})$ that is not far from \mathbf{y} . One elaboration that we leave to future work is whether the call to the dynamic combinatorial algorithm can be halted early as soon as the solution is provably different from \mathbf{y} .

Third, many set membership queries are avoided altogether. This can be done whenever the intersection of

$F^{-1}(\mathbf{y})$ with the subspace in which we are resampling is a convex set. This is clearly the case when $F^{-1}(\mathbf{y})$ is itself a convex set. Even in cases where $F^{-1}(\mathbf{y})$ is not a convex set, it may be that the intersection with the resampling subspace is convex. We can avoid certain set membership calls by the following reasoning: we are guaranteed inductively that the initial point \mathbf{w} lies within $F^{-1}(\mathbf{w})$. The alternative points that we encounter during slice sampling all lie on the line $\mathbf{w} + \alpha \cdot \mathbf{s}$ for some $\alpha \in \mathbb{R}$. Thus, if we observe that $\mathbf{w} + \alpha \cdot \mathbf{s}$ also lies within $F^{-1}(\mathbf{w})$, then we can infer that $\mathbf{w} + \alpha' \cdot \mathbf{s}$ lies within $F^{-1}(\mathbf{w})$ for all values $\alpha' \in [\alpha, 0]$ (if $\alpha < 0$) or $\alpha' \in [0, \alpha]$ (if $\alpha > 0$). When we encounter a point that can be proven to lie within $F^{-1}(\mathbf{w})$, there is clearly no reason to issue a set membership query.

When $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$ is log concave and $F^{-1}(\mathbf{w})$ is a convex set, then the restriction of $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$ to $F^{-1}(\mathbf{w})$ will also be log concave. Practically, this means that the slice sampler will never have to traverse difficult energy barriers to reach any part of the space. From a theoretical point of view, it guarantees that the slice sampler is “perfect” in that each step is drawing a uniform sample from the slice and is not simply leaving the uniform distribution invariant.

4.2 Hard EM

Rather than integrating $\int p(\mathbf{w} | \mathbf{y}, \boldsymbol{\psi}, \phi(\mathbf{x})) d\mathbf{w}$ in the E-step, we can instead maximize over \mathbf{w} . This yields a “hard EM” algorithm, in which the E-step seeks the \mathbf{w} in $F^{-1}(\mathbf{y})$ with maximum density $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$:

$$\operatorname{argmax}_{\mathbf{w}} p(\mathbf{w} | \mathbf{y}, \boldsymbol{\psi}, \phi(\mathbf{x})) = \operatorname{argmax}_{\mathbf{w} \in F^{-1}(\mathbf{y})} p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x})).$$

The methods presented in this section assume that $f_{\mathbf{w}}$ is defined as the exponential family form of Eq. 4, so $F^{-1}(\mathbf{y})$ is a convex polytope. If $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$ is log-concave in \mathbf{w} , the E-step is a convex optimization problem, constrained to a convex polytope. In the case that $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$ is Gaussian distributed, the optimization problem is a quadratic program (QP):

$$\min_{\mathbf{w}} \quad (\mathbf{w} - \hat{\mathbf{w}})^T (\mathbf{w} - \hat{\mathbf{w}}) \quad (13)$$

$$\text{s.t.} \quad f_{\mathbf{w}}(\mathbf{y}) \leq f_{\mathbf{w}'}(\mathbf{y}') \quad \forall \mathbf{y}' \neq \mathbf{y}, \quad (14)$$

where $\hat{\mathbf{w}}$ is the mode of $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$.

The difficulty is that the number of constraints needed to define $F^{-1}(\mathbf{y})$ (and thus the QP) is often exponential in $|\mathbf{y}|$. For all but small problems, it is intractable to explicitly instantiate the QP. However, as an efficient separation algorithm exists, the problem can be optimized by a cutting planes approach.

From an algorithmic perspective, we have arrived at an algorithm that bears some resemblance to a standard structural SVM formulation (Tsochantaridis et al.,

Algorithm 1 Hard EM+Margin: E-step

```

 $\mathbf{w} \leftarrow \hat{\mathbf{w}} \quad \mathcal{S} \leftarrow \emptyset \quad // \text{Active set of constraints.}$ 
while 1 do
   $\mathbf{y}' \leftarrow \operatorname{argmin}_{\hat{\mathbf{y}} \neq \mathbf{y}} f_{\mathbf{w}}(\hat{\mathbf{y}})$ 
  if  $\mathbf{y}' \in \mathcal{S}$  then
    break
  else
     $\mathcal{S} \leftarrow \mathcal{S} \cup \{\mathbf{y}'\}$ 
  end if
   $\mathbf{w} \leftarrow \operatorname{argmin}_{\mathbf{w}', \xi \geq 0} (\mathbf{w}' - \hat{\mathbf{w}})^T (\mathbf{w}' - \hat{\mathbf{w}}) - C\xi$ 
  s.t.  $f_{\mathbf{w}'}(\mathbf{y}) \leq f_{\mathbf{w}'}(\mathbf{y}') - \xi \quad \forall \mathbf{y}' \in \mathcal{S}$ 
end while
    
```

2005). The most obvious similarity is that both approaches involve formulating a QP that optimizes over model parameters, where the QP constraints are of the form that the ground truth \mathbf{y} has better objective than all other assignments \mathbf{y}' . To find violated constraints in the QP formulation of structural SVMs (or to find subgradients in the empirical risk minimization formulation) the same inner loop is used as in this hard EM formulation. The main difference in the inner loop is that the objective of the hard EM QP encourages \mathbf{w}' to be near $\hat{\mathbf{w}}$, the mode of $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$, versus the QP objective that encourages that \mathbf{w}' be near $\mathbf{0}$.

The second major difference is that the QP variables in this hard EM formulation are latent variables \mathbf{w} , while the variables in the QP from the structural SVM formulation are more analogous to our $\boldsymbol{\psi}$ parameters. In fact, we could make the connection to the structural SVM formulation clearer by making the additional assumption that the relationship between $\boldsymbol{\psi}$ and \mathbf{w} is deterministically set to $\mathbf{w}_n = \boldsymbol{\psi}^T \phi(\mathbf{x}_n)$. In this case, the QPs for all data cases n must be combined into one QP and solved simultaneously, like in a structural SVM algorithm. Mediating the interaction of $\boldsymbol{\psi}$ and \mathbf{y} via random \mathbf{w} allows for the probabilistic nature of the RandOM approach.

A practical problem with this hard EM algorithm is that the solutions \mathbf{w} often lie on the boundary of the polytope $F^{-1}(\mathbf{y})$. This leads to undesirable behavior in the M-step, where significant mass is placed outside $F^{-1}(\mathbf{y})$. To counteract this bias, we modify the QP objective to encourage finding points that have high density under $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$ but that are deeper inside $F^{-1}(\mathbf{y})$. We do this by adding a margin term to the QP objective, leading to the modified Hard EM QP given in Alg. 1. There, C is a parameter that trades off the margin between \mathbf{w} and the boundary of $F^{-1}(\mathbf{y})$ with maximizing $p(\mathbf{w} | \boldsymbol{\psi}, \phi(\mathbf{x}))$.

5 Experimental Evaluation

The main experimental questions we focus on concern the representational capabilities of Randomized Optimum models. We ran two sets of experiments. First, we experimented with synthetic data, and second, we applied a bipartite matching RandOM to a real-world lung CT scan registration task from medical imaging.


The goal is to learn to align landmarks in 3D lung volumes between two phases of the respiratory process. We show that RandOMs achieve higher accuracy than not only the Perturb-and-MAP models, but also a max-margin model trained structural SVM on the same features.

Evaluation: To evaluate the probability assigned to a data point (\mathbf{x}, \mathbf{y}) by the model, we draw samples \mathbf{y}_k from $p(\mathbf{y} | \psi, \phi(\mathbf{x}))$ using learned parameters ψ , and we count the fraction of samples that match the ground truth, i.e. $\frac{1}{K} \sum_k \mathbf{1}_{\{\mathbf{y}_k = \mathbf{y}_n\}}$. For our experiments here, we use $K=10000$.

Synthetic Matching Experiments: We generated synthetic data sets for bipartite matching models by sampling J landmark points (set \mathbf{A}) uniformly in the 2D box $[-1, 1]^2$. For each point, we also generated a real-valued size and 3D color vector in RGB space. We then generated J target points (set \mathbf{B}) by associating one point in \mathbf{B} with one point in \mathbf{A} , then drawing the properties of \mathbf{B} by perturbing the associated point in \mathbf{A} . In Dataset 1, we color the points in \mathbf{B} to be uniformly random (so color becomes an irrelevant feature), while noisily drawing size and position from a distribution centered at the value for the associated point in \mathbf{A} . In Dataset 2, we use a multimodal generation process; with probability $\frac{1}{2}$, we follow a similar procedure as above, but only set position to be relevant, and with probability $\frac{1}{2}$ set only color to be relevant. The features $\{\phi_{fij}(\mathbf{x}) | f \in \{\text{position, color, size}\}\}$ are length- J^2 real vectors, where there is an entry for each $(i \in \mathbf{A}, j \in \mathbf{B})$ pair. We set the entry of $\phi_{fij}(\mathbf{x})$ corresponding to (i, j) to be the squared distance between a_i and b_j along dimension f .

Results for the first data set are shown in Fig. 2. PM denotes Perturb-and-MAP, HEM denotes Hard EM, MCEM denotes Monte Carlo EM, and MM denotes max-margin. In parenthesis are the parameter setting, denoting the margin parameter, or the number of samples per iteration per instance. Comparing the MCEM algorithms to PM, we see that MCEM slightly outperforms PM, but that both do a good job in assigning mass to all training points. This is interesting, because it shows that Randomized Optimum models that do not attempt to approximate the standard Gibbs distribution can still have good representational power.

Quantitative results for both synthetic data sets are shown in Fig. 3. We compare the probability assigned to the data by each algorithm with several parameter settings. If the probability assigned to an instance was small enough that no samples were drawn that exactly matched the ground truth, then we remove that instance’s contribution from the log probability



(a) “Unimodal” data set. (b) “Multimodal” dataset.

Figure 2: Distribution of probability mass placed on the data by different learning algorithms. Along the x -axis are different training examples. The height of the bar shows the fraction of samples drawn from the model that matched the ground truth. $J = 5$, and $N = 10$.

(otherwise it would go to $-\infty$) and report this failure in the second column. Plots of how the log probability is assigned to data instances are shown in Fig. 2.

Amongst the Hard EM (HEM) algorithms, we see that including the term that rewards increasing the margin does improve performance. Increasing the term to be very large encourages finding the point that has largest margin, regardless of where mass of the distribution $p(\mathbf{w} | \psi, \phi(\mathbf{x}))$ lies. This causes the model to place nonzero mass on every instance, but it must increase its uncertainty, so the amount of mass placed on each instance decreases. We believe the reason why some data points are given negligible mass is because the E step is not guaranteed to find a point inside $F^{-1}(\mathbf{y})$ if the cutting planes algorithm is only run for a small number of iterations (as is done in these experiments).

As is particularly evident in the “multimodal” setting, the max-margin approach cannot always find a setting of ψ that perfectly separates the ground truths from all other assignments. In geometric terms, the subspaces spanned by the rows of $\phi(\mathbf{x}_n)$ do not intersect $F^{-1}(\mathbf{y}_n)$ for all n . While the probabilistic algorithms are able to place mass in all inverse mapping sets using distributions in \mathbf{w} space, the MM approach is not.

In both experiments, PM and MCEM achieve the best performance. Interestingly, while increasing the number of samples per iteration seems to improve MCEM, PM does not seem to benefit. In both cases, the top performing algorithm is MCEM with 100 samples drawn per instance per iteration in the E step.

Lung CT Registration Task: We obtained volumetric CT scans of pairs of lungs at different phases of the respiratory process, from 10 subjects (Castillo et al., 2009).¹ Included with the data are a set of hand-labeled landmarks giving ground truth correspondences between points in the pairs of lung vol-

¹Available at <http://www.dir-lab.com/index.html>

Alg	Avg. Log Prob	# Prob=0
Max-margin	0	2
PM(1)	-0.51	0
PM(10)	-0.50	0
PM(50)	-0.50	0
HEM(0.0)	-0.74	2
HEM(1.0)	-1.09	0
HEM(1e5)	-1.59	0
MCEM(1)	-0.44	0
MCEM(100)	-0.44	0

(a) “Unimodal” data

Alg	Avg. Log Prob	# Prob=0
Max-margin	0	4
PM(1)	-1.06	0
PM(10)	-1.06	0
PM(50)	-1.06	0
HEM(0.0)	-1.08	2
HEM(1.0)	-1.09	1
HEM(1e5)	-1.38	0
MCEM(1)	-1.07	0
MCEM(100)	-0.97	0

(b) “Multimodal” data

Figure 3: Quantitative results on synthetic data.

umes. We take 30 corresponding landmarks from each pair of lungs, and we extract several features for each landmark in each volume. We use a total of 20 features, including 3D position, mean intensity of varying sized patches around the landmark, standard deviation of intensity in varying sized patches around the landmark, and the Hessian-based “vesselness” filter of Frangi et al. (1998) at several scales. As in the synthetic data, for each landmark feature, we compute $\phi_{fij}(\mathbf{x})$ for a pair (a_i, b_j) as the squared difference of the features for a_i and b_j along dimension f .

We split the 10 instances into 10 different training sets, where we train on 9 instances and evaluate based on predictive performance on the held out 10th instance. The results are averaged over the 10 splits and are reported in Fig. 4. The columns show the fraction of predicted pairs that are correct (the “score”), along with average probability assigned to the test data, and average log probability assigned to the test data. Parentheses in the last column give the number of cases where the model assigned low enough probability to the test data such that we drew no samples that exactly matched the true test labeling.

We compare MCEM to PM and max-margin baselines. Here, though MCEM displays a small amount of overconfidence, it outperforms both baselines in overall score; that is, the Randomized Optimum models produce the most accurate predictions on this task.

6 Discussion and Related Work

To our knowledge, the line of work that includes Papandreou and Yuille (2011) and this work is unique in

Alg	Avg. Score	Avg. Prob	Avg. Log Prob
Max-margin	91.0%	.400	0 (6)
PM(1)	95.0%	.538	-1.52 (0)
PM(10)	94.9%	.543	-1.54 (0)
MCEM(1)	98.8%	.837	-1.67 (2)
MCEM(5)	96.8%	.568	-0.62 (1)
MCEM(10)	96.7%	.544	-0.56 (1)

Figure 4: Quantitative results on lung CT data.

the technique of defining generative probabilistic models that include an explicit optimization procedure. However, the motivating ideas are present in several works, which build probabilistic models around specialized tractable computational structures.

For example, Determinantal Point Processes (Kulesza and Taskar, 2011) define probability models around the computation of the determinant of a matrix. Domke (2011) defines probability models around a fixed number of iterations of belief propagation. Neither of these computations is a discrete optimization of the type we consider, but the spirits are similar.

There is also work on learning under the assumption that test-time inference will involve running a discrete optimization procedure. Structured output SVMs are perhaps the most popular, and they have been widely applied e.g., Taskar et al. (2005); Ratliff et al. (2006); Szummer et al. (2008); Joachims et al. (2009).

In this paper we have presented a framework for defining probabilistic models that exploit efficient methods for optimizing the scoring function. The models derive from a simple generative procedure that allows a rich class of probabilistic models. While we have shown how a previous model, Perturb-and-MAP, can be applied beyond settings typically modeled as conditional random fields, RandOMs are more general than even the extended version of Perturb-and-MAP. The learning procedures we provide show how to directly formulate a maximum likelihood learning objective. The experimental results show this learning to be both tractable and successful in a real-world application.

RandOMs trade enumerating an exponential number of configurations for an integration of the inverse mapping polytope. We believe moving the difficult computations to the continuous latent space is desirable, because it allows the use of powerful continuous-variable tools to be used in the context of structured discrete output spaces.

Finally, we have only begun to explore the representational power of the framework. RandOMs can readily incorporate other generative distributions over the latent variables, such as Gaussian Processes, mixture models, or various nonparametric models. Exploring these alternatives is a focus of our future work.

Acknowledgements

We thank Juan Eugenio Iglesias for help designing features for the lungs model, along with help preparing the data.

References

- R. P. Adams, I. Murray, and D. J. C. MacKay. The Gaussian process density sampler. In *NIPS 21*, 2009.
- R. Castillo, E. Castillo, R. Guerra, V. Johnson, T. McPhail, A. Garg, and T. Guerrero. A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. *Phys Med Biol*, 54:1849–1870, 2009.
- W. H. Cunningham. Minimum cuts, modular functions, and matroid polyhedra. *Networks*, 15(2):205–215, 1985.
- J. Domke. Parameter learning with truncated message-passing. In *CVPR*, 2011.
- A. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *MICCAI*, 1998.
- L. A. Hannah, D. M. Blei, and W. B. Powell. Dirichlet process mixtures of generalized linear models. *JMLR*, 12:1923–1953, 2011.
- T. Joachims, T. Finley, and C.-N. Yu. Cutting-plane training of structural SVMs. *Machine Learning*, 77(1):27–59, 2009.
- V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE PAMI*, 26(2):147–159, 2004.
- G. Korsah, A. Stentz, and M. Dias. The dynamic Hungarian algorithm for the assignment problem with changing costs. Technical Report CMU-RI-TR-07-27, Robotics Institute, Pittsburgh, PA, July 2007.
- A. Kulesza and B. Taskar. Structured determinantal point processes. In *NIPS 23*, 2011.
- Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. Huang. A tutorial on energy-based learning. In *Predicting Structured Data*. MIT Press, 2006.
- J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, 1957.
- I. Murray, Z. Ghahramani, and D. J. C. MacKay. MCMC for doubly-intractable distributions. In *UAI*, pages 359–366, 2006.
- R. M. Neal. Slice sampling. *Annals of Statistics*, 31(3):705–767, 2003.
- V. Ng. Graph-cut-based anaphoricity determination for coreference resolution. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2009.
- G. Papandreou and A. Yuille. Perturb-and-MAP random fields: Using discrete optimization to learn and sample from energy models. In *ICCV*, 2011.
- N. Ratliff, J. A. Bagnell, and M. Zinkevich. Maximum margin planning. In *ICML*, 2006.
- C. Rother, V. Kolmogorov, and A. Blake. Grab-Cut: interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics*, 23(3):309–314, 2004.
- C. Smith. A characterization of star-shaped sets. *American Mathematical Monthly*, 75(4), 1968.
- M. Szummer, P. Kohli, and D. Hoiem. Learning CRFs using graph cuts. In *ECCV*, 2008.
- B. Taskar, S. Lacoste-Julien, and D. Klein. A discriminative matching approach to word alignment. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 73–80, 2005.
- L. Torresani, V. Kolmogorov, and C. Rother. Feature correspondence via graph matching: Models and global optimization. In *ECCV*. 2008.
- I. Tsochantaris, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *JMLR*, 6:1453–1484, 2005.
- S. C. Turaga, K. L. Briggman, M. Helmstaedter, W. Denk, and H. S. Seung. Maximin affinity learning of image segmentation. In *NIPS*, 2010.
- L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979.
- Wainwright and Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 2008.
- H. Wang. *Proteins, Interactions, and Complexes: A Computational Approach*. PhD thesis, Department of Computer Science, Stanford University, 2008.
- G. C. G. Wei and M. A. Tanner. A Monte Carlo implementation of the EM algorithm and the poor man’s data augmentation algorithms. *Journal of the American Statistical Association*, 85(411):pp. 699–704, 1990.