
Discriminative Embeddings of Latent Variable Models for Structured Data

Hanjun Dai, Bo Dai
Le Song

{HANJUNDAI, BODAI}@GATECH.EDU
LSONG@CC.GATECH.EDU

College of Computing, Georgia Institute of Technology, Atlanta, USA

Abstract

Kernel classifiers and regressors designed for structured data, such as sequences, trees and graphs, have significantly advanced a number of interdisciplinary areas such as computational biology and drug design. Typically, kernels are designed beforehand for a data type which either exploit statistics of the structures or make use of probabilistic generative models, and then a discriminative classifier is learned based on the kernels via convex optimization. However, such an elegant two-stage approach also limited kernel methods from scaling up to millions of data points, and exploiting discriminative information to learn feature representations.

We propose, `structure2vec`, an effective and scalable approach for structured data representation based on the idea of embedding latent variable models into feature spaces, and learning such feature spaces using discriminative information. Interestingly, `structure2vec` extracts features by performing a sequence of function mappings in a way similar to graphical model inference procedures, such as mean field and belief propagation. In applications involving millions of data points, we showed that `structure2vec` runs 2 times faster, produces models which are 10,000 times smaller, while at the same time achieving the state-of-the-art predictive performance.

1. Introduction

Structured data, such as sequences, trees and graphs, are prevalent in a number of interdisciplinary areas such as protein design, genomic sequence analysis, and drug design (Schölkopf et al., 2004). To learn from such complex data, we have to first transform such data explicitly or implicitly into some vectorial representations,

and then apply machine learning algorithms in the resulting vector space. So far kernel methods have emerged as one of the most effective tools for dealing with structured data, and have achieved the state-of-the-art classification and regression results in many sequence (Leslie et al., 2002a; Vishwanathan & Smola, 2003) and graph datasets (Gärtner et al., 2003; Borgwardt, 2007).

The success of kernel methods on structured data relies crucially on the design of kernel functions — positive semidefinite similarity measures between pairs of data points (Schölkopf & Smola, 2002). By designing a kernel function, we have implicitly chosen a corresponding feature representation for each data point which can potentially has infinite dimensions. Later learning algorithms for various tasks and with potentially very different nature can then work exclusively on these pairwise kernel values without the need to access the original data points. Such modular structure of kernel methods has been very powerful, making them the most elegant and convenient methods to deal with structured data. Thus designing kernel for different structured objects, such as strings, trees and graphs, has always been an important subject in the kernel community. However, in the big data era, this modular framework has also limited kernel methods in terms of their ability to scale up to millions of data points, and exploit discriminative information to learn feature representations.

For instance, a class of kernels are designed based on the idea of “bag of structures” (BOS), where each structured data point is represented as a vector of counts for elementary structures. The spectrum kernel and variants for strings (Leslie et al., 2002a), subtree kernel (Ramon & Gärtner, 2003), graphlet kernel (Shervashidze et al., 2009) and Weisfeiler-lehman graph kernel (Shervashidze et al., 2011) all follow this design principle. In other words, the feature representations of these kernels are fixed before learning, with each dimension corresponding to a substructure, independent of the supervised learning tasks at hand. Since there are many unique substructures which may or may not be useful for the learning tasks, the explicit feature space of such kernels typically has very high dimensions. Subsequently algorithms dealing with

the pairwise kernel values have to work with a big kernel matrix squared in the number of data points. The square dependency on the number of data points largely limits these BOS kernels to datasets of size just thousands.

A second class of kernels are based on the ingenious idea of exploiting the ability of probabilistic graphical models (GM) in describing noisy and structured data to design kernels. For instance, one can use hidden Markov models for sequence data, and use pairwise Markov random fields for graph data. The Fisher kernel (Jaakkola & Haussler, 1999) and probability product kernel (Jebara et al., 2004) are two representative instances within the family. The former method first fits a common generative model to the entire dataset, and then uses the empirical Fisher information matrix and the Fisher score of each data point to define the kernel; The latter method instead fits a different generative model for each data point, and then uses inner products between distributions to define the kernel. Typically the parameterization of these GM kernels are chosen before hand. Although the process of fitting generative models allow the kernels to adapt to the geometry of the input data, the resulting feature representations are still independent of the discriminative task at hand. Furthermore, the extra step of fitting generative models to data can be a challenging computation and estimation task by itself, especially in the presence of latent variables. Very often in practice, one finds that BOS kernels are easier to deploy than GM kernels, although the latter is supposed to capture the additional geometry and uncertainty information of data.

In this paper, we wish to revisit the idea of using graphical models for kernel or feature space design, with the goal of scaling up kernel methods for structured data to millions of data points, and allowing the kernel to learn the feature representation from label information. Our idea is to model each structured data point as a latent variable model, then embed the graphical model into feature spaces (Smola et al., 2007; Song et al., 2009), and use inner product in the embedding space to define kernels. Instead of fixing a feature or embedding space beforehand, we will also learn the feature space by directly minimizing the empirical loss defined by the label information.

The resulting embedding algorithm, `structure2vec`, runs in a scheme similar to graphical model inference procedures, such as mean field and belief propagation. Instead of performing probabilistic operations (such as sum, product and renormalization), the algorithm performs nonlinear function mappings in each step, inspired by kernel message passing algorithm in Song et al. (2010; 2011). Furthermore, `structure2vec` is also different from the kernel message passing algorithm in several aspects. First, `structure2vec` deals with a different

scenario, *i.e.*, learning similarity measure for structured data. Second, `structure2vec` learns the nonlinear mappings using the discriminative information. And third, a variant of `structure2vec` can run in a mean field update fashion, different from message passing algorithms.

Besides the above novel aspects, `structure2vec` is also very scalable in terms of both memory and computation requirements. First, it uses a small and explicit feature map for the nonlinear feature space, and avoids the need for keeping the kernel matrix. This makes the subsequent classifiers or regressors order of magnitude smaller compared to other methods. Second, the nonlinear function mapping in `structure2vec` can be learned using stochastic gradient descent, allowing it to handle extremely large scale datasets.

Finally in experiments, we show that `structure2vec` compares favorably to other kernel methods in terms of classification accuracy in medium scale sequence and graph benchmark datasets including SCOP and NCI. Furthermore, `structure2vec` can handle extremely large data set, such as the 2.3 million molecule dataset from Harvard Clean Energy Project, run 2 times faster, produce model 10,000 times smaller and achieve state-of-the-art accuracy. These strong empirical results suggest that the graphical models, theoretically well-grounded methods for capturing structure in data, combined with embedding techniques and discriminative training can significantly improve the performance in many large scale real-world structured data classification and regression problems.

2. Backgrounds

We denote by X a random variable with domain \mathcal{X} , and refer to instantiations of X by the lower case character, x . We denote a density on \mathcal{X} by $p(X)$, and denote the space of all such densities by \mathcal{P} . We will also deal with multiple random variables, X_1, X_2, \dots, X_ℓ , with joint density $p(X_1, X_2, \dots, X_\ell)$. For simplicity of notation, we assume that the domains of all $X_t, t \in [\ell]$ are the same, but the methodology applies to the cases where they have different domains. In the case when \mathcal{X} is a discrete domain, the density notation should be interpreted as probability, and integral should be interpreted as summation instead. Furthermore, we denote by H a hidden variable with domain \mathcal{H} and distribution $p(H)$. We use similar notation convention for variable H and X .

Kernel Methods. Suppose the structured data is represented by $\chi \in \mathcal{G}$. Kernel methods owe the name to the use of kernel functions, $k(\chi, \chi') : \mathcal{G} \times \mathcal{G} \mapsto \mathbb{R}$, which are symmetric positive semidefinite (PSD), meaning that for all $n > 1$, and $\chi_1, \dots, \chi_n \in \mathcal{G}$, and $c_1, \dots, c_n \in \mathbb{R}$, we have $\sum_{i,j=1}^n c_i c_j k(\chi_i, \chi_j) \geq 0$. A signature of kernel methods is that learning algorithms for various tasks and with potentially very different nature can work exclusively

on these pairwise kernel values without the need to access the original data points.

Kernels for Structured Data. Each kernel function will correspond to some feature map $\phi(\chi)$, where the kernel function can be expressed as the inner product between feature maps, *i.e.*, $k(\chi, \chi') = \langle \phi(\chi), \phi(\chi') \rangle$. For structured input domain, one can design kernels using counts on substructures. For instance, the spectrum kernel for two sequences χ and χ' is defined as (Leslie et al., 2002a)

$$k(\chi, \chi') = \sum_{s \in \mathcal{S}} \#(s \in \chi) \#(s \in \chi') \quad (1)$$

where \mathcal{S} is the set of possible subsequences, $\#(s \in x)$ counts the number occurrence of subsequence s in x . In this case, the feature map $\phi(\chi) = (\#(s_1 \in \chi), \#(s_2 \in \chi), \dots)^\top$ corresponds to a vector of dimension $|\mathcal{S}|$. Similarly, the graphlet kernel (Shervashidze et al., 2009) for two graphs χ and χ' can also be defined as (1), but \mathcal{S} is now the set of possible subgraphs, and $\#(s \in \chi)$ counts the number occurrence of subgraphs. We refer to this class of kernels as ‘‘bag of structures’’ (BOS) kernel.

Kernels can also be defined by leveraging the power of probabilistic graphical models. For instance, the Fisher kernel (Jaakkola & Haussler, 1999) is defined using a parametric model $p(\chi|\theta^*)$ around its maximum likelihood estimate θ^* , *i.e.*, $k(\chi, \chi') = U_\chi^\top I^{-1} U_{\chi'}$, where $U_\chi := \nabla_{\theta=\theta^*} \log p(\chi|\theta)$ and $I = \mathbb{E}_{\mathcal{G}}[U_{\mathcal{G}} U_{\mathcal{G}}^\top]$ is the Fisher information matrix. Another classical example along the line is the probability product kernel (Jebara et al., 2004). Different from the Fisher kernel based on generative model fitted with the whole dataset, the probability product kernel is calculated based on the models $p(\chi|\theta)$ fitted to individual data point, *i.e.*, $k(\chi, \chi') = \int_{\mathcal{G}} p(\tau|\theta_\chi)^\rho p(\tau|\theta_{\chi'})^\rho d\tau$ where θ_χ and $\theta_{\chi'}$ are the maximum likelihood parameters for data point χ and χ' respectively. We refer to this class of kernels as the ‘‘graphical model’’ (GM) kernels.

Hilbert Space Embedding of Distributions. Hilbert space embeddings of distributions are mappings of distributions into potentially *infinite* dimensional feature spaces (Smola et al., 2007),

$$\mu_X := \mathbb{E}_X [\phi(X)] = \int_{\mathcal{X}} \phi(x) p(x) dx : \mathcal{P} \mapsto \mathcal{F} \quad (2)$$

where the distribution is mapped to its expected feature map, *i.e.*, to a point in a feature space. Kernel embedding of distributions has rich representational power. Some feature map can make the mapping injective (Sriperumbudur et al., 2008), meaning that if two distributions, $p(X)$ and $q(X)$, are different, they are mapped to two distinct points in the feature space. For instance, when $\mathcal{X} = \mathbb{R}^d$, the feature spaces of many commonly used kernels, such as the Gaussian RBF kernel $\exp(-\|x - x'\|_2^2)$, can make the embedding injective.

Alternatively, one can treat an injective embedding μ_X of

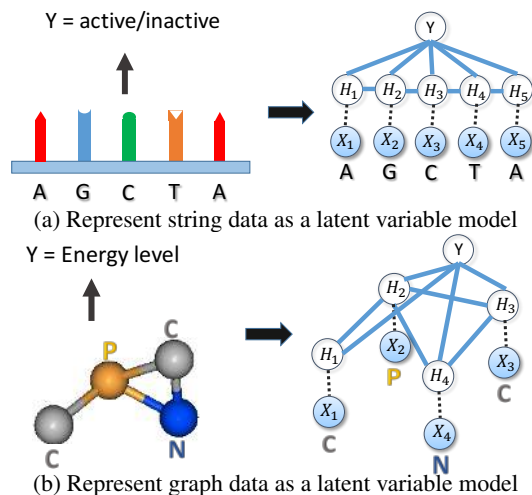


Figure 1. Building graphical model with hidden variables from structured string and general graph data. Y is the supervised information, which can be real number (for regression) or discrete integer (for classification).

a density $p(X)$ as a sufficient statistic of the density. Any information we need from the density is preserved in μ_X : with μ_X one can uniquely recover $p(X)$, and any operation on $p(X)$ can be carried out via a corresponding operation on μ_X with the same result. For instance, this property will allow us to compute a functional $f : \mathcal{P} \mapsto \mathbb{R}$ of the density using the embedding only, *i.e.*,

$$f(p(x)) = \tilde{f}(\mu_X) \quad (3)$$

where $\tilde{f} : \mathcal{F} \mapsto \mathbb{R}$ is a corresponding function applied on μ_X . Similarly the property can also be generalized to operators. For instance, applying an operator $\mathcal{T} : \mathcal{P} \mapsto \mathbb{R}^d$ to a density can also be equivalently carried out using its embedding, *i.e.*,

$$\mathcal{T} \circ p(x) = \tilde{\mathcal{T}} \circ \mu_X, \quad (4)$$

where $\tilde{\mathcal{T}} : \mathcal{F} \mapsto \mathbb{R}^d$ is the alternative operator working on the embedding. In our later sections, we will extensively exploit this property of injective embeddings, by assuming that there exists a feature space such that the embeddings are injective. We include the discussion of other related work in Appendix A.

3. Model for a Structured Data Point

Without loss of generality, we assume each structured data point χ is a graph, with a set of nodes $\mathcal{V} = \{1, \dots, V\}$ and a set of edges \mathcal{E} . We will use x_i to denote the value of the attribute for node i . We note the node attributes are different from the label of the entire data point. For instance, each atom in a molecule will correspond to a node in the graph, and the node attribute will be the atomic number, while the label for the entire molecule can be whether the molecule is a good drug or not. Other structures, such as sequences and trees, can be viewed as special cases of general graphs.

We will model the structured data point χ as an instance

drawn from a graphical model. More specifically, we will model the label of each node in the graph with a variable X_i , and furthermore, associate an additional hidden variable H_i with it. Then we will define a pairwise Markov random field on these collection of random variables

$$p(\{H_i\}, \{X_i\}) \propto \prod_{i \in \mathcal{V}} \Phi(H_i, X_i) \prod_{(i,j) \in \mathcal{E}} \Psi(H_i, H_j) \quad (5)$$

where Ψ and Φ are nonnegative node and edge potentials respectively. In this model, the variables are connected according to the graph structure of the input data point. That is to say, we use the graph structure of the input data directly as the conditional independence structure of an undirected graphical model. Figure 1 illustrates two concrete examples in constructing the graphical models for strings and graphs. One can design more complicated graphical models which go beyond pairwise Markov random fields, and consider longer range interactions with potentials involving more variables. We will focus on pairwise Markov random fields for simplicity of representation.

We note that such a graphical model is built for each individual data point, and the conditional independence structures of two graphical models can be different if the two data points χ and χ' are different. Furthermore, we do not observe the value for the hidden variables $\{H_i\}$, which makes the learning of the graphical model potentials Φ and Ψ even more difficult. Thus, we will not pursue the standard route of maximum likelihood estimation, and rather we will consider the sequence of computations needed when we try to embed the posterior of $\{H_i\}$ into a feature space.

4. Embedding Latent Variable Models

We will embed the posterior marginal $p(H_i | \{x_i\})$ of a hidden variable using a feature map $\phi(H_i)$, *i.e.*,

$$\mu_i = \int_{\mathcal{H}} \phi(h_i) p(h_i | \{x_i\}) dh_i. \quad (6)$$

The exact form of $\phi(H_i)$ and the parameters in MRF $p(H_i | \{x_i\})$ is not fixed at the moment, and we will learn them later using supervision signals for the ultimate discriminative target. For now, we will assume that $\phi(H_i) \in \mathbb{R}^d$ is a finite dimensional feature space, and the exact value of d will be determined by cross-validation in later experiments. However, computing the embedding is a very challenging task for general graphs: it involves performing an inference in graphical model where we need to integrate out all variables except H_i , *i.e.*,

$$p(H_i | \{x_i\}) = \int_{\mathcal{H}^{\mathcal{V}-1}} p(H_i, \{h_j\} | \{x_j\}) \prod_{j \in \mathcal{V} \setminus i} dh_j. \quad (7)$$

Only when the graph structure is a tree, exact computation can be carried out efficiently via message passing (Pearl, 1988). Thus in the general case, approximate inference

algorithms, *e.g.*, mean field inference and loopy belief propagation (BP), are developed. In many applications, however, these variational inference algorithms exhibit excellent empirical performance (Murphy et al., 1999). Several theoretical studies have also provided insight into the approximations made by loopy BP, partially justifying its application to graphs with cycles (Wainwright & Jordan, 2008; Yedidia et al., 2001a).

In the following subsection, we will explain the embedding of mean field and loopy BP. The embedding of other variational inference methods, *e.g.*, double-loop BP, damped BP, tree-reweighted BP, and generalized BP will be explained in Appendix C. We show that the iterative update steps in these algorithms, which are essentially minimizing approximations to the exact free energy, can be simply viewed as function mappings of the embedded marginals using the alternative view in (3) and (4).

4.1. Embedding Mean-Field Inference

The vanilla mean-field inference tries to approximate $p(\{H_i\} | \{x_i\})$ with a product of *independent* density components $p(\{H_i\} | \{x_i\}) \approx \prod_{i \in \mathcal{V}} q_i(h_i)$ where each $q_i(h_i) \geq 0$ is a valid density, such that $\int_{\mathcal{H}} q_i(h_i) dh_i = 1$. Furthermore, these density components are found by minimizing the following variational free energy (Wainwright & Jordan, 2008),

$$\min_{q_1, \dots, q_d} \int_{\mathcal{H}^d} \prod_{i \in \mathcal{V}} q_i(h_i) \log \frac{\prod_{i \in \mathcal{V}} q_i(h_i)}{p(\{h_i\} | \{x_i\})} \prod_{i \in \mathcal{V}} dh_i.$$

One can show that the solution to the above optimization problem needs to satisfy the following fixed point equations for all $i \in \mathcal{V}$

$$\log q_i(h_i) = c_i + \log(\Phi(h_i, x_i)) + \quad (8)$$

$$\sum_{j \in \mathcal{N}(i)} \int_{\mathcal{H}} q_j(h_j) \log(\Psi(h_i, h_j) \Phi(h_j, x_j)) dh_j$$

where $\mathcal{N}(i)$ are the set of neighbors of variable H_i in the graphical model, and c_i is a constant. The fixed point equations in (8) imply that $q_i(h_i)$ is a functional of a set of neighboring marginals $\{q_j\}_{j \in \mathcal{N}(i)}$, *i.e.*,

$$q_i(h_i) = f\left(h_i, x_i, \{q_j\}_{j \in \mathcal{N}(i)}, \{x_j\}_{j \in \mathcal{N}(i)}\right). \quad (9)$$

If for each marginal q_i , we have an injective embedding

$$\tilde{\mu}_i = \int_{\mathcal{H}} \phi(h_i) q_i(h_i) dh_i,$$

then, using similar reasoning as in (3), we can equivalently express the fixed point equation from an embedding point of view, *i.e.*, $q_i(h_i) = \tilde{f}(h_i, x_i, \{\tilde{\mu}_j\}_{j \in \mathcal{N}(i)}, \{x_j\}_{j \in \mathcal{N}(i)})$, and consequently using the operator view from (4), we have

$$\tilde{\mu}_i = \tilde{\mathcal{T}} \circ \left(x_i, \{\tilde{\mu}_j\}_{j \in \mathcal{N}(i)}, \{x_j\}_{j \in \mathcal{N}(i)}\right). \quad (10)$$

For the embedded mean field (10), the function \tilde{f} and operator $\tilde{\mathcal{T}}$ have complicated nonlinear dependencies on the potential functions Ψ , Φ , and the feature mapping

Algorithm 1 Embedding Mean Field

```

1: Input: parameter  $\mathbf{W}$  in  $\tilde{\mathcal{T}}$ 
2: Initialize  $\tilde{\mu}_i^{(0)} = \mathbf{0}$ , for all  $i \in \mathcal{V}$ 
3: for  $t = 1$  to  $T$  do
4:   for  $i \in \mathcal{V}$  do
5:      $l_i = \sum_{j \in \mathcal{N}(i)} \tilde{\mu}_i^{(t-1)}$ 
6:      $\tilde{\mu}_i^{(t)} = \sigma(W_1 x_i + W_2 l_i + W_3 \sum_{j \in \mathcal{N}(i)} x_j)$ 
7:   end for
8: end for {fixed point equation update}
9: return  $\{\tilde{\mu}_i^T\}_{i \in \mathcal{V}}$ 
    
```

ϕ which is unknown and need to be learned from data. Instead of first learning the Ψ and Φ , and then working out $\tilde{\mathcal{T}}$, we will pursue a different route where we directly parameterize $\tilde{\mathcal{T}}$ and later learn it with supervision signals.

In terms of the parameterization, we will assume $\tilde{\mu}_i \in \mathbb{R}^d$ where d is a hyperparameter chosen using cross-validation. For $\tilde{\mathcal{T}}$, one can use any nonlinear function mappings. For instance, we can parameterize it as a neural network

$$\tilde{\mu}_i = \sigma\left(W_1 x_i + W_2 \sum_{j \in \mathcal{N}(i)} \tilde{\mu}_j + W_3 \sum_{j \in \mathcal{N}(i)} x_j\right) \quad (11)$$

where $\sigma(\cdot) := \max\{0, \cdot\}$ is a rectified linear unit applied elementwisely to its argument, and $\mathbf{W} = \{W_1, W_2, W_3\}$. The number of the rows in \mathbf{W} equals to d . With such parameterization, the mean field iterative update in the embedding space can be carried out as Algorithm 1. We could also multiply $\tilde{\mu}_i$ with V to rescale the range of message embeddings if needed. In fact, with or without V , the functions will be the same in terms of the representation power. Specifically, for any (\mathbf{W}, V) , we can always find another ‘equivalent’ parameters (\mathbf{W}', I) where $\mathbf{W}' = \{W_1, W_2 V, W_3\}$.

4.2. Embedding Loopy Belief Propagation

Loopy belief propagation is another variational inference method, which essentially optimizes the Bethe free energy taking *pairwise* interactions into account (Yedidia et al., 2001b),

$$\min_{\{q_{ij}\}_{(i,j) \in \mathcal{E}}} - \sum_i (|\mathcal{N}(i)| - 1) \int_{\mathcal{H}} q_i(h_i) \log \frac{q_i(h_i)}{\Phi(h_i, x_i)} dh_i + \sum_{i,j} \int_{\mathcal{H}^2} q_{ij}(h_i, h_j) \log \frac{q_{ij}(h_i, h_j)}{\Psi(h_i, h_j) \Phi(h_i, x_i) \Phi(h_j, x_j)} dh_i dh_j$$

subject to pairwise marginal consistency constraints: $\int_{\mathcal{H}} q_{ij}(h_i, h_j) dh_j = q_i(h_i)$, $\int_{\mathcal{H}} q_{ij}(h_i, h_j) dh_i = q_j(h_j)$, and $\int_{\mathcal{H}} q_i(h_i) dh_i = 1$. One can obtain the fixed point condition for the above optimization for all $(i, j) \in \mathcal{E}$,

$$m_{ij}(h_j) \propto \int_{\mathcal{H}} \prod_{k \in \mathcal{N}(i) \setminus j} m_{ki}(h_i) \Phi_i(h_i, x_i) \Psi_{ij}(h_i, h_j) dh_i, \quad (12)$$

$$q_i(h_i) \propto \Phi(h_i, x_i) \prod_{j \in \mathcal{N}(i)} m_{ji}(h_i).$$

where $m_{ij}(h_j)$ is the intermediate result called the message from node i to j . Furthermore, $m_{ij}(h_j)$ is a nonnegative

Algorithm 2 Embedding Loopy BP

```

1: Input: parameter  $\mathbf{W}$  in  $\tilde{\mathcal{T}}_1$  and  $\tilde{\mathcal{T}}_2$ 
2: Initialize  $\tilde{\nu}_{ij}^{(0)} = \mathbf{0}$ , for all  $(i, j) \in \mathcal{E}$ 
3: for  $t = 1$  to  $T$  do
4:   for  $(i, j) \in \mathcal{E}$  do
5:      $\tilde{\nu}_{ij}^t = \sigma(W_1 x_i + W_2 \sum_{k \in \mathcal{N}(i) \setminus j} \tilde{\nu}_{ki}^{(t-1)})$ 
6:   end for
7: end for
8: for  $i \in \mathcal{V}$  do
9:    $\tilde{\mu}_i = \sigma(W_3 x_i + W_4 \sum_{k \in \mathcal{N}(i)} \tilde{\nu}_{ki}^{(T)})$ 
10: end for
11: return  $\{\tilde{\mu}_i\}_{i \in \mathcal{V}}$ 
    
```

function which can be normalized to a density, and hence can also be embedded.

Similar to the reasoning in the mean field case, the (12) implies the messages $m_{ij}(h_j)$ and marginals $q_i(h_i)$ are functionals of messages from neighbors, *i.e.*,

$$m_{ij}(h_j) = f(h_j, x_i, \{m_{ki}\}_{k \in \mathcal{N}(i) \setminus j}),$$

$$q_i(h_i) = g(h_i, x_i, \{m_{ki}\}_{k \in \mathcal{N}(i)}).$$

With the assumption that there is an injective embedding for each message $\tilde{\nu}_{ij} = \int \phi(h_j) m_{ij}(h_j) dh_j$ and for each marginal $\tilde{\mu}_i = \int \phi(h_i) q_i(h_i) dh_i$, we can apply the reasoning from (3) and (4), and express the messages and marginals from the embedding view,

$$\tilde{\nu}_{ij} = \tilde{\mathcal{T}}_1 \circ (x_i, \{\tilde{\nu}_{ki}\}_{k \in \mathcal{N}(i) \setminus j}), \quad (13)$$

$$\tilde{\mu}_i = \tilde{\mathcal{T}}_2 \circ (x_i, \{\tilde{\nu}_{ki}\}_{k \in \mathcal{N}(i)}). \quad (14)$$

We will also use parametrization for loopy BP embedding similar to the mean field case, *i.e.*, neural network with rectified linear unit σ . Specifically, assume $\tilde{\nu}_{ij} \in \mathbb{R}^d$, $\tilde{\mu}_i \in \mathbb{R}^d$

$$\tilde{\nu}_{ij} = \sigma\left(W_1 x_i + W_2 \sum_{k \in \mathcal{N}(i) \setminus j} \tilde{\nu}_{ki}\right) \quad (15)$$

$$\tilde{\mu}_i = \sigma\left(W_3 x_i + W_4 \sum_{k \in \mathcal{N}(i)} \tilde{\nu}_{ki}\right) \quad (16)$$

where $\mathbf{W} = \{W_1, W_2, W_3, W_4\}$ are matrices with appropriate sizes. Note that one can use other nonlinear function mappings to parameterize $\tilde{\mathcal{T}}_1$ and $\tilde{\mathcal{T}}_2$ as well. Overall, the loopy BP embedding updates is summarized in Algorithm 2.

With similar strategy as in mean field case, we will learn the parameters in $\tilde{\mathcal{T}}_1$ and $\tilde{\mathcal{T}}_2$ later with supervision signals from the discriminative task.

4.3. Embedding Other Variational Inference

In fact, there are many other variational inference methods, with different forms of free energies or different optimization algorithms, resulting different message update forms, *e.g.*, double-loop BP (Yuille, 2002), damped BP (Minka, 2001), tree-reweightd BP (Wainwright et al.,

2003), and generalized BP (Yedidia et al., 2001b). The proposed embedding method is a general technique which can be tailored to these algorithms. The major difference is the dependences in the messages. For the details of embedding of these algorithms, please refer to Appendix C.

5. Discriminative Training

Similar to kernel BP (Song et al., 2010; 2011) and kernel EP (Jitkrittum et al., 2015), our current work exploits feature space embedding to reformulate graphical model inference procedures. However, different from the kernel BP and kernel EP, in which the feature spaces are chosen beforehand and the conditional embedding operators are learned locally, our approach will learn both the feature spaces, the transformation \tilde{T} , as well as the regressor or classifier for the target values end-to-end using label information.

Specifically, we are provided with a training dataset $\mathcal{D} = \{\chi_n, y_n\}_{n=1}^N$, where χ_n is a structured data point and $y_n \in \mathcal{Y}$, where $\mathcal{Y} = \mathbb{R}$ for regression or $\mathcal{Y} = \{1, \dots, K\}$ for classification problem, respectively. With the feature embedding procedure introduced in Section 4, each data point will be represented as a set of embeddings $\{\tilde{\mu}_i^n\}_{i \in V_n} \in \mathcal{F}$. Now the goal is to learn a regression or classification function f linking $\{\tilde{\mu}_i^n\}_{i \in V_n}$ to y_n .

More specifically, in the case of regression problem, we will parametrize function $f(\chi_n)$ as $u^\top \sigma(\sum_{i=1}^{V_n} \tilde{\mu}_i^n)$, where $u \in \mathbb{R}^d$ is the final mapping from summed (or pooled) embeddings to output. The parameters u and those \mathbf{W} involved in the embeddings are learned by minimizing the empirical square loss

$$\min_{u, \mathbf{W}} \sum_{n=1}^N \left(y_n - u^\top \sigma \left(\sum_{i=1}^{V_n} \tilde{\mu}_i^n \right) \right)^2. \quad (17)$$

Note that each data point will have its own graphical model and embedded features due to its individual structure, but the parameters u and \mathbf{W} , are shared across these graphical models.

In the case of K -class classification problem, we denote z is the 1-of- K representation of y , i.e., $z \in \{0, 1\}^K$, $z^k = 1$ if $y = k$, and $z^i = 0, \forall i \neq k$. By adopt the softmax loss, we obtain the optimization for embedding parameters and discriminative classifier estimation as,

$$\min_{u = \{u^k\}_{k=1}^K, \mathbf{W}} \sum_n \sum_{k=1}^K -z_n^k \log u^k \sigma \left(\sum_{i=1}^{V_n} \tilde{\mu}_i^n \right), \quad (18)$$

where $u = \{u^k\}_{k=1}^K$, $u^k \in \mathbb{R}^d$ are the parameters for mapping embedding to output.

The same idea can also be generalized to other discriminative tasks with different loss functions. As we can see from the optimization problems (17) and (18), the objective functions are directly related to the corresponding discriminative tasks, and so as to \mathbf{W} and u .

Algorithm 3 Discriminative Embedding

Input: Dataset $\mathcal{D} = \{\chi_n, y_n\}_{n=1}^N$, loss function $l(f(\chi), y)$.

Initialize $U^0 = \{\mathbf{W}^0, \mathbf{u}^0\}$ randomly.

for $t = 1$ **to** T **do**

Sample $\{\chi_t, y_t\}$ uniform randomly from \mathcal{D} .

Construct latent variable model $p(\{H_i^t\}|\chi_n)$ as (5).

Embed $p(\{H_i^t\}|\chi_n)$ as $\{\tilde{\mu}_i^n\}_{i \in V_n}$ by Algorithm 1 or 2 with \mathbf{W}^{t-1} .

Update $U^t = U^{t-1} + \lambda_t \nabla_{U^{t-1}} l(f(\tilde{\mu}^n; U^{t-1}), y_n)$.

end for

return $U^T = \{\mathbf{W}^T, \mathbf{u}^T\}$

Conceptually, the procedure starts with representing each datum by a graphical model constructed corresponding to its *individual* structure with *sharing* potential functions, and then, we embed these graphical models with the *same* feature mappings. Finally the embedded marginals are aggregated with a prediction function for a discriminative task. The shared potential functions, feature mappings and final prediction functions are all learned together for the ultimate task with supervision signals.

We optimize the objective (17) or (18) with stochastic gradient descent for scalability consideration. However, other optimization algorithms are also applicable, and our method does not depend on this particular choice. The gradients of the parameters \mathbf{W} are calculated recursively similar to recurrent neural network for sequence models. In our case, the recursive structure will correspond the message passing structure. The overall framework is illustrated in Algorithm 3. For details of the gradient calculation, please refer to Appendix D.

6. Experiments

Below we first compare our method with algorithms using prefixed kernel on string and graph benchmark datasets. Then we focus on Harvard Clean Energy Project dataset which contains 2.3 million samples. We demonstrate that while getting comparable performance on medium sized datasets, we are able to handle millions of samples, and getting much better when more training data are given. The two variants of `structure2vec` are denoted as DE-MF and DE-LBP, which stands for discriminative embedding using mean field or loopy belief propagation, respectively.

Our algorithms are implemented with C++ and CUDA, and experiments are carried out on clusters equipped with NVIDIA Tesla K20. The code is available on <https://github.com/Hanjun-Dai/graphnn>.

6.1. Benchmark structure datasets

We compare our algorithm on string benchmark datasets with the kernel method with existing sequence kernels, i.e., the spectrum string kernel (Leslie et al., 2002a), mismatch string kernel (Leslie et al., 2002b) and fisher kernel with

HMM generative models (Jaakkola & Haussler, 1999). On graph benchmark datasets, we compare with subtree kernel (Ramon & Gärtner, 2003) (R&G, for short), random walk kernel (Gärtner et al., 2003; Vishwanathan et al., 2010), shortest path kernel (Borgwardt & Kriegel, 2005), graphlet kernel (Shervashidze et al., 2009) and the family of Weisfeiler-Lehman kernels (WL kernel) (Shervashidze et al., 2011). After getting the kernel matrix, we train SVM classifier or regressor on top.

Without explicitly mentioned, we perform cross validation for all methods, and report the average performance. We include the details of tuning hyper parameters for baselines and our methods in Appendix E.2.

6.1.1. STRING DATASET

Here we do experiments on two string binary classification benchmark datasets. The first one (denoted as SCOP) contains 7329 sequences obtained from SCOP (Structural Classification of Proteins) 1.59 database (Andreeva et al., 2004). Methods are evaluated on the ability to detect members of a target SCOP family (positive test set) belonging to the same SCOP superfamily as the positive training sequences, and no members of the target family are available during training. We use the same 54 target families and the same training/test splits as in remote homology detection (Kuang et al., 2005). The second one is FC and RES dataset (denoted as FC_RES) provided by CRISPR/Cas9 system, on which the task is to tell whether the guide RNA will direct Cas9 to target DNA. There are 5310 guides included in the dataset. Details of this dataset can be found in Doench et al. (2014); Fusi et al. (2015). We use two variants for spectrum string kernel: 1) kmer-single, where the constructed kernel matrix $K_k^{(s)}$ only consider patterns of length k ; 2) kmer-concat, where kernel matrix $K^{(c)} = \sum_{i=1}^k K_k^{(s)}$. We also find the normalized kernel matrix $K_k^{Norm}(x, y) = \frac{K_k(x, y)}{\sqrt{K_k(x, x)K_k(y, y)}}$ helps.

| | FC_RES | SCOP |
|-------------|----------------------|----------------------|
| kmer-single | 0.7606±0.0187 | 0.7097±0.0504 |
| kmer-concat | 0.7576±0.0235 | 0.8467±0.0489 |
| mismatch | 0.7690±0.0197 | 0.8637±0.1192 |
| fisher | 0.7332±0.0314 | 0.8662±0.0879 |
| DE-MF | 0.7713±0.0208 | 0.9068±0.0685 |
| DE-LBP | 0.7701±0.0225 | 0.9167±0.0639 |

Table 1. Mean AUC on string classification datasets

Table 1 reports the mean AUC of different algorithms. We found two variants of `structure2vec` are consistently better than the string kernels. Also, the improvement in SCOP is more significant than in FC_RES. This is because SCOP is a protein dataset and its alphabet size $|\Sigma|$ is much larger than that of FC_RES, an RNA dataset. Furthermore, the dimension of the explicit features for a k-mer kernel is $O(|\Sigma|^k)$, which can make the off-diagonal entries of

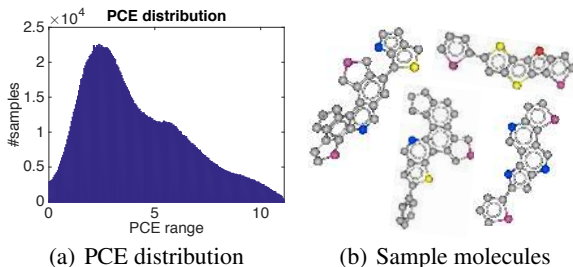


Figure 3. PCE value distribution and sample molecules from CEP dataset. Hydrogens are not displayed.

kernel matrix very small (or even zero) with large alphabet size and k . That’s also the reason why kmer-concat performs better than kmer-single. `structure2vec` learns a discriminative feature space, rather than prefix it beforehand, and hence does not have this problem.

6.1.2. GRAPH DATASETS

We test the algorithms on five benchmark datasets for graph kernel: MUTAG, NCI1, NCI109, ENZYMES and D&D. MUTAG (Debnath et al., 1991). NCI1 and NCI109 (Wale et al., 2008) are chemical compounds dataset, while ENZYMES (Borgwardt & Kriegel, 2005) and D&D (Dobson & Doig, 2003) are of proteins. The task is to do multi-class or binary classification. For more details of dataset, please refer to Appendix E.1.

The results of baseline algorithms are taken from Shervashidze et al. (2011) since we use exactly the same setting here. From the accuracy comparison shown in Figure 2, we can see the proposed embedding methods are comparable to the alternative graph kernels, on different graphs with different number of labels, nodes and edges. Also, in dataset D&D which consists of 82 different types of labels, our algorithm performs much better. As reported in Shervashidze et al. (2011), the time required for constructing dictionary for the graph kernel can take up to more than a year of CPU time in this dataset, while our algorithm can learn the discriminative embedding efficiently from structured data directly without the construction of the handcraft dictionary.

6.2. Harvard Clean Energy Project(CEP) dataset

The Harvard Clean Energy Project (Hachmann et al., 2011) is a theory-driven search for the next generation of organic solar cell materials. One of the most important properties of molecule for this task is the overall efficiency of the energy conversion process in a solar cell, which is determined by the power conversion efficiency (PCE). The Clean Energy Project (CEP) performed expensive simulations for the 2.3 million candidate molecules on IBMs World Community Grid, in order to get this property value. So using machine learning approach to accurately predict the PCE values is a promising direction for the high throughput screening and discovering new materials.

In this experiment, we randomly select 90% of the data for




Figure 2. 10-fold cross-validation accuracies on graph classification benchmark datasets. The ‘sp’ in the figure stands for shortest-path.

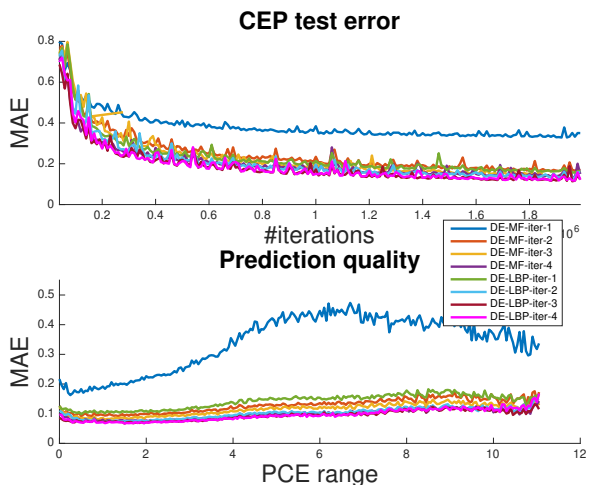


Figure 4. Details of training and prediction results for DE-MF and DE-LBP with different number of fixed point iterations.

training, and the rest 10% for testing. This setting is similar to Pyzer-Knapp et al. (2015), except that we use the entire 2.3m dataset here. Since the data is distributed unevenly (see Figure 3), we resampled the training data (but not the test data) to make the algorithm put more emphasis on molecules with higher PCE values, in order to make accurate prediction for promising candidate molecules. Since the traditional kernel methods are not scalable, we make the explicit feature maps for WL subtree kernel by collecting all the molecules and creating dictionary for the feature space. The other graph kernels, like edge kernel and shortest path kernel, are having too large feature dictionary to work with. We use RDKit (Landrum, 2012) to extract features for atoms (nodes) and bonds (edges).

The mean absolute error (MAE) and root mean square error (RMSE) are reported in Table 2. We found utilizing graph information can accurately predict PCE values. Also, our proposed two methods are working equally well. Although WL tree kernel with degree 6 is also working well, it requires 10,000 times more parameters than *structure2vec* and runs 2 times slower. The preprocessing needed for WL tree kernel also makes it difficult to use in large datasets.

To understand the effect of the inference embedding in the proposed algorithm framework, we further compare our methods with different number of fixed point iterations

| | test MAE | test RMSE | # params |
|----------------|---------------|---------------|----------|
| Mean Predictor | 1.9864 | 2.4062 | 1 |
| WL lv-3 | 0.1431 | 0.2040 | 1.6m |
| WL lv-6 | 0.0962 | 0.1367 | 1378m |
| DE-MF | 0.0914 | 0.1250 | 0.1m |
| DE-LBP | 0.0850 | 0.1174 | 0.1m |

Table 2. Test prediction performance on CEP dataset. WL lv- k stands for Weisfeiler-lehman with degree k .

in Figure 4. It can see that, higher number of fixed point iterations will lead to faster convergence, though the number of parameters of the model in different settings are the same. The mean field embedding will get much worse result if only one iteration is executed. Compare to the loopy BP case with same setting, the latter one will always have one more round message passing since we need to aggregate the messages from edge to node in the last step. And also, from the quality of prediction we find that, though making slightly higher prediction error for molecules with high PCE values due to insufficient data, the variants of our algorithm are not overfitting the ‘easy’ (i.e., the most popular) range of PCE value.

7. Conclusion

We propose, *structure2vec*, an effective and scalable approach for structured data representation based on the idea of embedding latent variable models into feature spaces, and learning such feature spaces using discriminative information. Interestingly, *structure2vec* extracts features by performing a sequence of function mappings in a way similar to graphical model inference procedures, such as mean field and belief propagation. In applications involving millions of data points, we showed that *structure2vec* runs 2 times faster, produces models 10,000 times smaller, while at the same time achieving the state-of-the-art predictive performance. *structure2vec* provides a nice example for the general strategy of combining the strength of graphical models, Hilbert space embedding of distribution and deep learning approach, which we believe will become common in many other learning tasks.

Acknowledgements. This project was supported in part by NSF/NIH BIGDATA 1R01GM108341, ONR N00014-15-1-2340, NSF IIS-1218749, and NSF CAREER IIS-1350983.

References

- Andreeva, A., Howorth, D., Brenner, S. E., Hubbard, T. J., Chothia, C., and Murzin, A. G. Scop database in 2004: refinements integrate structure and sequence family data. *Nucleic acids research*, 32(suppl 1):D226–D229, 2004.
- Borgwardt, K. M. *Graph Kernels*. PhD thesis, Ludwig-Maximilians-University, Munich, Germany, 2007.
- Borgwardt, K. M. and Kriegel, H.-P. Shortest-path kernels on graphs. In *ICDM*, 2005.
- Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- Chang, C. C. and Lin, C. J. *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Chen, L. C., Schwing, A. G., Yuille, A. L., and Urtasun, R. Learning deep structured models. *arXiv preprint arXiv:1407.2538*, 2014.
- Debnath, A. K., Lopez de Compadre, R. L., Debnath, G., Shusterman, A. J., and Hansch, C. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *J Med Chem*, 34:786–797, 1991.
- Dobson, P. D. and Doig, A. J. Distinguishing enzyme structures from non-enzymes without alignments. *J Mol Biol*, 330(4): 771–783, Jul 2003.
- Doench, J. G., Hartenian, E., Graham, D. B., Tothova, Z., Hegde, H., Smith, I., Sullender, M., Ebert, B. L., Xavier, R. J., and Root, D. E. Rational design of highly active sgRNAs for crispr-cas9-mediated gene inactivation. *Nature biotechnology*, 32(12):1262–1267, 2014.
- Duvenaud, D. K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, pp. 2215–2223, 2015.
- Fusi, N., Smith, I., Doench, J., and Listgarten, J. In silico predictive modeling of crispr/cas9 guide efficiency. *bioRxiv*, 2015. doi: 10.1101/021568. URL <http://biorxiv.org/content/early/2015/06/26/021568>.
- Gärtner, T., Flach, P.A., and Wrobel, S. On graph kernels: Hardness results and efficient alternatives. In Schölkopf, B. and Warmuth, M. K. (eds.), *Proceedings of Annual Conference. Computational Learning Theory*, pp. 129–143. Springer, 2003.
- Caruana, R., Lawrence, S., and Giles, L. Overfitting in neural nets: Backpropagation, conjugate gradient, and early stopping. In *Advances in Neural Information Processing Systems 13*, volume 13, pp. 402. MIT Press, 2001.
- Hachmann, J., Olivares-Amaya, R., Atahan-Evrenk, S., Amador-Bedolla, C., Sánchez-Carrera, R. S., Gold-Parker, A., Vogt, L., Brockway, A. M., and Aspuru-Guzik, A. The harvard clean energy project: large-scale computational screening and design of organic photovoltaics on the world community grid. *The Journal of Physical Chemistry Letters*, 2(17):2241–2251, 2011.
- Henaff, M., Bruna, J., and LeCun, Y. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.
- Hershey, J. R., Roux, J. L., and Wenginger, F. Deep unfolding: Model-based inspiration of novel deep architectures. *arXiv preprint arXiv:1409.2574*, 2014.
- Heskes, T. Stable fixed points of loopy belief propagation are local minima of the bethe free energy. *Advances in Neural Information Processing Systems*, pp. 343–350. MIT Press, 2002.
- Jaakkola, T. S. and Haussler, D. Exploiting generative models in discriminative classifiers. In Kearns, M. S., Solla, S. A., and Cohn, D. A. (eds.), *Advances in Neural Information Processing Systems 11*, pp. 487–493. MIT Press, 1999.
- Jebara, T., Kondor, R., and Howard, A. Probability product kernels. *J. Mach. Learn. Res.*, 5:819–844, 2004.
- Jitkrittum, W., Gretton, A., Heess, N., Eslami, S. M. A., Lakshminarayanan, B., Sejdinovic, D., and Szabó, Z. Kernel-based just-in-time learning for passing expectation propagation messages. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence, UAI 2015, July 12-16, 2015, Amsterdam, The Netherlands*, pp. 405–414, 2015.
- Kuang, R., Ie, E., Wang, K., Wang, K., Siddiqi, M., Freund, Y., and Leslie, C. Profile-based string kernels for remote homology detection and motif extraction. *Journal of bioinformatics and computational biology*, 3(03):527–550, 2005.
- Landrum, G. Rdkit: Open-source cheminformatics (2013), 2012.
- Leslie, C., Eskin, E., and Noble, W. S. The spectrum kernel: A string kernel for SVM protein classification. In *Proceedings of the Pacific Symposium on Biocomputing*, pp. 564–575. Singapore, 2002a. World Scientific Publishing.
- Leslie, C., Eskin, E., Weston, J., and Noble, W. S. Mismatch string kernels for SVM protein classification. In *Advances in Neural Information Processing Systems*, volume 15, Cambridge, MA, 2002b. MIT Press.
- Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R.. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- Lin, G., Shen, C., Reid, I., and van den Hengel, A. Deeply learning the messages in message passing inference. In *Advances in Neural Information Processing Systems*, 2015.
- Minka, T. The EP energy function and minimization schemes. See www.stat.cmu.edu/minka/papers/learning.html, August, 2001.
- Mou, L., Li, G., Zhang, L., Wang, T., and Jin, Z.. Convolutional neural networks over tree structures for programming language processing. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- Murphy, K. P., Weiss, Y., and Jordan, M. I. Loopy belief propagation for approximate inference: An empirical study. In *UAI*, pp. 467–475, 1999.

- Pearl, J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, 1988.
- Pyzer-Knapp, E. O., Li, K., and Aspuru-Guzik, A. Learning from the harvard clean energy project: The use of neural networks to accelerate materials discovery. *Advanced Functional Materials*, 25(41):6495–6502, 2015.
- Ramon, J. and Gärtner, T. Expressivity versus efficiency of graph kernels. Technical report, First International Workshop on Mining Graphs, Trees and Sequences (held with ECML/PKDD’03), 2003.
- Ross, S., Munoz, D., Hebert, M., and Bagnell, J. A. Learning message-passing inference machines for structured prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2737–2744. IEEE, 2011.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- Schölkopf, B., Tsuda, K., and Vert, J.-P. *Kernel Methods in Computational Biology*. MIT Press, Cambridge, MA, 2004.
- Schölkopf, B., and Smola, A. J. *Learning with Kernels*. MIT Press, Cambridge, MA, 2002.
- Shervashidze, N., Vishwanathan, S. V. N., Petri, T., Mehlhorn, K., and Borgwardt, K. Efficient graphlet kernels for large graph comparison. *Proceedings of International Conference on Artificial Intelligence and Statistics*. Society for Artificial Intelligence and Statistics, 2009.
- Shervashidze, N., Schweitzer, P., Van Leeuwen, E. J., Mehlhorn, K., and Borgwardt, K. M. Weisfeiler-lehman graph kernels. *The Journal of Machine Learning Research*, 12:2539–2561, 2011.
- Smola, A. J., Gretton, A., Song, L., and Schölkopf, B. A Hilbert space embedding for distributions. In *Proceedings of the International Conference on Algorithmic Learning Theory*, volume 4754, pp. 13–31. Springer, 2007.
- Song, L., Huang, J., Smola, A. J., and Fukumizu, K. Hilbert space embeddings of conditional distributions. In *Proceedings of the International Conference on Machine Learning*, 2009.
- Song, L., Gretton, A., and Guestrin, C. Nonparametric tree graphical models. In *13th Workshop on Artificial Intelligence and Statistics*, volume 9 of *JMLR workshop and conference proceedings*, pp. 765–772, 2010.
- Song, L., Gretton, A., Bickson, D., Low, Y., and Guestrin, C. Kernel belief propagation. In *Proc. Intl. Conference on Artificial Intelligence and Statistics*, volume 10 of *JMLR workshop and conference proceedings*, 2011.
- Sriperumbudur, B., Gretton, A., Fukumizu, K., Lanckriet, G., and Schölkopf, B. Injective Hilbert space embeddings of probability measures. In *Proceedings of Annual Conference. Computational Learning Theory*, pp. 111–122, 2008.
- Sugiyama, M. and Borgwardt, K. Halting in random walk kernels. In *Advances in Neural Information Processing Systems*, pp. 1630–1638, 2015.
- Vishwanathan, S. V. N. and Smola, A. J. Fast kernels for string and tree matching. In Becker, S., Thrun, S., and Obermayer, K. (eds.), *Advances in Neural Information Processing Systems 15*, pp. 569–576. MIT Press, Cambridge, MA, 2003.
- Vishwanathan, S. V. N., Schraudolph, N. N., Kondor, I. R., and Borgwardt, K. M. Graph kernels. *Journal of Machine Learning Research*, 2010. URL <http://www.stat.purdue.edu/~vishy/papers/VisSchKonBor10.pdf>. In press.
- Wainwright, M., Jaakkola, T., and Willsky, A. Tree-reweighted belief propagation and approximate ML estimation by pseudo-moment matching. In *9th Workshop on Artificial Intelligence and Statistics*, 2003.
- Wainwright, M. J. and Jordan, M. I. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.
- Wale, N., Watson, I. A., and Karypis, G. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3):347–375, 2008.
- Yedidia, J. S., Freeman, W. T., and Weiss, Y. Generalized belief propagation. *Advances in Neural Information Processing Systems*, pp. 689–695. MIT Press, 2001a.
- Yedidia, J.S., Freeman, W.T., and Weiss, Y. Bethe free energy, kikuchi approximations and belief propagation algorithms. Technical report, Mitsubishi Electric Research Laboratories, 2001b.
- Yedidia, J.S., Freeman, W.T., and Weiss, Y. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51(7): 2282–2312, 2005.
- Yuille, A. L. Cccp algorithms to minimize the bethe and kikuchi free energies: Convergent alternatives to belief propagation. *Neural Computation*, 14:2002, 2002.
- Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, B., Su, Z., Du, D., Huang, C., and Torr, P. Conditional random fields as recurrent neural networks. *arXiv preprint arXiv:1502.03240*, 2015.