
NEURAL MESSAGE PASSING FOR MULTI-LABEL CLASSIFICATION

Jack Lanchantin, Arshdeep Sekhon, & Yanjun Qi

Department of Computer Science
University of Virginia
Charlottesville VA 22903
{jj15sw,as5cu,yq2h}@virginia.edu

ABSTRACT

Multi-label classification (MLC) is the task of assigning a set of target labels for a given sample. Modeling the combinatorial label interactions in MLC has been a long-haul challenge. We propose Label Message Passing (LaMP) Neural Networks to efficiently model the joint prediction of multiple labels. LaMP treats labels as nodes on a label-interaction graph and computes the hidden representation of each label node conditioned on the input using attention-based neural message passing. Attention enables LaMP to assign different importances to neighbor nodes per label, learning how labels interact (implicitly). The proposed models are simple, accurate, interpretable, structure-agnostic, and applicable for predicting dense labels since LaMP is incredibly parallelizable. We validate the benefits of LaMP on seven real-world MLC datasets, covering a broad spectrum of input/output types and outperforming the state-of-the-art results. Notably, LaMP enables intuitive interpretation of how classifying each label depends on the elements of a sample and at the same time rely on its interaction with other labels¹.

1 Introduction

Multi-label classification (MLC) is receiving increasing attention in areas such as natural language processing, computational biology, and image recognition. Accurate and scalable MLC methods are in urgent need for applications like assigning topics to web articles, or identifying binding proteins on DNA. The most common and straightforward MLC method is the binary relevance (BR) approach that considers multiple target labels independently [46]. However, in many MLC tasks there is a clear dependency structure among labels, which BR methods ignore.

Unfortunately, accurately modelling all combinatorial label interactions is an NP-hard problem. Many types of models, including a few deep neural network (DNN) based, have been introduced to approximately model such interactions, thus boosting classification accuracy.

Our main concern of this paper is how to represent multiple labels jointly (and conditioned on the input features) in order to make accurate predictions. The most relevant literature addressing this concern falls roughly into three groups.

The first group, probabilistic classifier chain (PCC) models, formulate the joint label dependencies using the chain rule and perform MLC in a sequential prediction manner [39, 55, 36]. Notably, [36] used a recurrent neural network (RNN) sequence to sequence (Seq2Seq) architecture [17] for MLC and achieved the state-of-the-art performance on multiple text-based datasets. However, these methods are inherently unfit for MLC tasks due to their incapacity to be parallelized, and inability to perform well in dense label settings, or when there are a large number of positive labels (since errors propagate in the sequential prediction). We refer the reader to the supplementary material for a full

¹We provide our code and datasets at <https://github.com/QData/LaMP>

background and analysis of PCC methods (Appendix section 5). The second group learns a shared latent space representing both input features and output labels, and then upsamples from the space to reconstruct the target labels [57, 6]. The main drawback of this group is the interpretability issue with a learned low dimensional latent space, as many real-world applications prefer interpretable predictors. The third group models conditional label dependencies using a structured output or graphical model representation [29, 45]. However, these methods are often limited to only considering pair-wise dependencies due to computational constraints, or are forced to use some variation of approximate inference which has no clear representation of conditional dependencies.

Thus our main question is: *is it possible to have accurate, flexible and explainable MLC methods that are applicable to many dense labels?* This paper provides empirical results showing that this is possible through extending attention based Message Passing Neural Networks (MPNNs) to learn the joint representation of multiple labels conditioned on input features.

MPNNs [15] are a class of methods that efficiently learn the joint representations of variables using neural message passing strategies. They provide a flexible framework for modeling multiple variables jointly which have no explicit ordering.

The key idea of our method is to rely on attention-based neural message passing entirely to draw global dependencies from labels to input features, and from labels to labels. To the best of our knowledge, this is the first extension of MPNNs to model a conditional joint representation of output labels, and additionally the first extension of MPNNs to model the interactions of variables where the exact structure is unknown. We name the proposed method Label Message Passing (LaMP) Networks since it performs neural message passing on an unknown, fully-connected label-to-label graph. Through intra-attention (aka self-attention), LaMP assigns different importance to different neighbor nodes per label, dynamically learning how labels interact conditioned on a specific input. We further extend LaMP to cases when a known label interaction graph is provided by modifying the intra-attention to only attend over a node’s known neighbors. LaMP networks allow for parallelization in training and testing and can work with dense labels, overcoming the drawbacks of PCC methods.

LaMP most closely belongs to the third MLC category we mentioned above, however it trains a unified model to classify each label *and* model the label to label dependencies at the same time, in an end-to-end fashion. The important aspect is that LaMP networks automatically learn the output label dependency structure conditioned on a specific input using neural message passing. This in turn can easily be interpreted to understand the conditional structure.

The main contributions of this paper include: (1) **Accurate MLC**: Our model achieves similar, or better performance compared to the previous state of the art across five MLC metrics. We validate our model on eight MLC datasets which cover a wide spectrum of input data structure: sequences (English text, DNA), tabular (binary word vectors), graph (drug molecules), and images, as well as output label structure: unknown and graph. (2) **Interpretable**: Although deep-learning based systems have widely been viewed as “black boxes”, our attention based LaMP models allow for a straightforward way to extract three different types of model visualization: intermediate network predictions, label to feature dependencies, and label to label dependencies.

2 Method: LaMP Networks

Notations. We define the following notations, used throughout the paper. Let $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ be the set of data samples with inputs $\mathbf{x} \in X$ and outputs $\mathbf{y} \in Y$. Inputs \mathbf{x} are a (possibly ordered) set of S components $\{x_1, x_2, \dots, x_S\}$, and outputs \mathbf{y} are a set of L labels $\{y_1, y_2, \dots, y_L\}$. MLC involves predicting the set of binary labels $\{y_1, y_2, \dots, y_L\}, y_i \in \{0, 1\}$ given input \mathbf{x} .

In general we can assume to represent the input feature components as embedded vectors $\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_S\}, \mathbf{z}_i \in \mathbb{R}^d$, using some learned embedding matrix $\mathbf{W}^x \in \mathbb{R}^{\delta \times d}$. Here d is the embedding size and, δ is the size of x_i . x_i can be any component of a particular input (for example, words in a sentence, patches of an image, nodes of a known graph, or one of the tabular features).

Similarly, labels can be first represented as embedded vectors $\{\mathbf{u}_1^{t=0}, \mathbf{u}_2^{t=0}, \dots, \mathbf{u}_L^{t=0}\}, \mathbf{u}_i^t \in \mathbb{R}^d$, through a learned embedding matrix $\mathbf{W}^y \in \mathbb{R}^{L \times d}$, where L denotes the number of labels. Here we use t to represent the ‘state’ of the embedding after the t^{th} update step. This is because in LaMP networks, each label embedding is updated for t steps before the predictions are made. The key idea of LaMP networks is that labels are represented as nodes in a label-interaction graph G_{yy} denoting

nodes as embedding vectors $\{\mathbf{u}_{1:L}^t\}$. LaMP networks use MPNN modules with attention to pass messages from input embeddings $\{\mathbf{z}_{1:S}\}$ to G_{yy} , and then within G_{yy} to model the joint prediction of labels.

2.1 Background: Message Passing Neural Networks

Message Passing Neural Networks (MPNNs) [15] are a generalization of graph neural networks (GNNs) [41]. MPNNs model variables as nodes on a graph G . Here $G = (V, E)$, where V describes the set of nodes (variables) and E denotes the set of edges (about how variables interact with other variables). In MPNNs, joint representations of nodes and edges are modelled using message passing rather than explicit probabilistic formulations, allowing for efficient inference. MPNNs model the joint dependencies using message function M^t and node update function U^t for T time steps, where t is the current time step. The hidden state $\mathbf{v}_i^t \in \mathbb{R}^d$ of node $i \in G$ is updated based on messages \mathbf{m}_i^t from its neighboring nodes $\{\mathbf{v}_{j \in \mathcal{N}(i)}^t\}$ defined by neighborhood $\mathcal{N}(i)$:

$$\mathbf{m}_i^t = \sum_{j \in \mathcal{N}(i)} M^t(\mathbf{v}_i^t, \mathbf{v}_j^t), \quad (1)$$

$$\mathbf{v}_i^{t+1} = U^t(\mathbf{m}_i^t). \quad (2)$$

After T rounds of iterative updates to spread information to distant nodes, a readout function R is used on the updated node embeddings to make predictions like classifying nodes or classifying properties about the graph.

Many possibilities exist for functions M^t and U^t . We specifically choose to pass messages using intra-attention (also called as self-attention) neural message passing which enable nodes to attend over their neighborhoods differentially. This allows for the network to learn different importances for different nodes in a neighborhood, without depending on knowing the graph structure upfront (essentially learning the unknown graph structure) [53]. In this formulation, messages for node \mathbf{v}_i^t are obtained by a weighted sum of all its neighboring nodes $\{\mathbf{v}_{j \in \mathcal{N}(i)}^t\}$ where the weights are calculated by attention representing the importance of each neighbor for a specific node [2]. In the rest of the paper, we use ‘‘graph attention’’ and ‘‘neural message passing’’ interchangeably.

Intra-attention neural message passing works as follows. We first calculate attention weights α_{ij}^t for pair of nodes $(\mathbf{v}_i^t, \mathbf{v}_j^t)$ using attention function $a(\cdot)$:

$$\alpha_{ij}^t = \text{softmax}_j(e_{ij}^t) = \frac{\exp(e_{ij}^t)}{\sum_{k \in \mathcal{N}(i)} \exp(e_{ik}^t)} \quad (3)$$

$$e_{ij}^t = a(\mathbf{v}_i^t, \mathbf{v}_j^t) \quad (4)$$

$$a(\mathbf{v}_i^t, \mathbf{v}_j^t) = \frac{(\mathbf{W}^q \mathbf{v}_i^t)^\top (\mathbf{W}^u \mathbf{v}_j^t)}{\sqrt{d}} \quad (5)$$

where e_{ij}^t represents the importance of node j for node i , however un-normalized. e_{ij}^t are normalized across all neighboring nodes of node i using a softmax function (Eq 3) to get α_{ij}^t . For the attention function $a(\cdot)$, we used a scaled dot product with node-wise linear transformations $\mathbf{W}^q \in \mathbb{R}^{d \times d}$ on node \mathbf{v}_i^t and $\mathbf{W}^u \in \mathbb{R}^{d \times d}$ on node \mathbf{v}_j^t . Scaling by \sqrt{d} is used to mitigate training issues [52].

Then we use a so called attention message function M_{atn}^t to produce the message from node j to node i using the learned attention weights α_{ij}^t and another transformation matrix $\mathbf{W}^v \in \mathbb{R}^{d \times d}$:

$$M_{\text{atn}}(\mathbf{v}_i^t, \mathbf{v}_j^t; \mathbf{W}) = \alpha_{ij}^t \mathbf{W}^v \mathbf{v}_j^t, \quad (6)$$

$$\mathbf{m}_i^t = \mathbf{v}_i^t + \sum_{j \in \mathcal{N}(i)} M_{\text{atn}}(\mathbf{v}_i^t, \mathbf{v}_j^t; \mathbf{W}). \quad (7)$$

Eq 7 computes the full message \mathbf{m}_i^t for node \mathbf{v}_i^t by linearly combining messages from all neighbor nodes $j \in \mathcal{N}(i)$ with a residual connection on the current \mathbf{v}_i^t .

Lastly, node v_i^t is updated to next state v_i^{t+1} using message m_i^t by a multi-layer perceptron (MLP) update function U_{mlp} , plus a m_i^t residual connection:

$$U_{\text{mlp}}(m_i^t; \mathbf{W}) = \text{ReLU}(\mathbf{W}^r m_i^t + b_1)^\top \mathbf{W}^b + b_2 \quad (8)$$

$$v_i^{t+1} = m_i^t + U_{\text{mlp}}(m_i^t; \mathbf{W}). \quad (9)$$

Function U_{mlp} is parameterized with matrices $\{\mathbf{W}^r \in \mathbb{R}^{d \times d}, \mathbf{W}^b \in \mathbb{R}^{d \times d}\}$. It is important to note that \mathbf{W} in Eq 9 are shared (i.e., separately applied) across all nodes. This can be viewed as 1-dimensional convolution operation with kernel and stride sizes of 1. Weight sharing across nodes is a key aspect of MPNNs, where node dependencies are learned in an order-invariant manner.

2.2 LaMP: Label Message Passing

Given the input embeddings $\{z_1, z_2, \dots, z_S\}$, the goal of Label Message Passing is to model the conditional dependencies between label embeddings $\{u_1^t, u_2^t, \dots, u_L^t\}$ using Message Passing Neural Networks. We assume that the label embeddings are nodes on a label-interaction graph called G_{yy} , where the initial state of the embeddings $\{u_{1:L}^0\}$ at $t = 0$ are obtained using label embedding matrix \mathbf{W}^y .

Each step t in Label Message Passing consists of two parts in order to update the label embeddings: (a). Feature-to-Label Message Passing, where messages are passed from the input embeddings to the label embeddings, and (b). Label-to-Label Message Passing, where messages are passed between labels. An overview of our model is shown in Fig. 1. We explain these two parts in detail in the following subsections. LaMP Networks use T steps of attention-based neural message passing to update the label nodes before a readout function makes a prediction for each label i on its final state u_i^T .

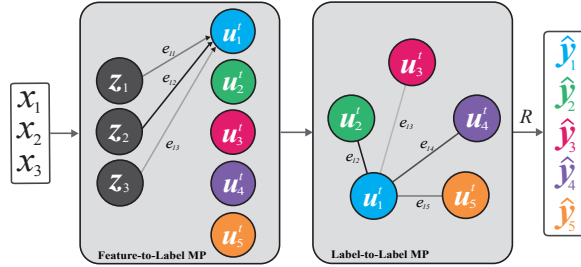


Figure 1: **LaMP Networks.** Given input \mathbf{x} , we encode its components $\{x_1, x_2, x_3\}$ as embedded input nodes $\{z_1, z_2, z_3\}$. We encode labels $\{y_1, y_2, \dots, y_5\}$ as embedded label nodes $\{u_1^t, u_2^t, \dots, u_5^t\}$ of label-interaction graph G_{yy} . First, MPNN_{xy} is used to pass messages from the input nodes to the labels nodes and update the label nodes. Then, MPNN_{yy} is used to pass messages between the label nodes and update label nodes. Finally, readout function R performs node-level classification on label nodes to make binary label predictions $\{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_5\}$.

Updating Label Embeddings via Feature-to-Label Message Passing

Given a particular input \mathbf{x} with embedded feature components $\{z_1, z_2, \dots, z_S\}$, the first step in LaMP is to update the label embeddings by passing messages from the input embeddings to the label embeddings, as shown in the ‘‘Feature-to-Label MP’’ block of Fig. 1. To do this, LaMP uses neural message passing module MPNN_{xy} to update the i^{th} label node’s embedding u_i^t using the embeddings of all the components of an input.

That is, we update each u_i^t by using a weighted sum of all input embeddings $\{z_{1:S}\}$, in which the weights represent how important an input component is to the i^{th} label node. The weights for the message are learned via Label-to-Feature attention (i.e., each label attends to each input embedding differently to compute the weights).

In this step, messages are only passed from the input nodes to the label nodes, and not vice versa (i.e. Feature-to-Label message passing is directed).

More specifically, to update label embedding u_i^t , MPNN_{xy} uses attention message function M_{atn}^t on all embeddings of the input $\{z_{1:S}\}$ to produce messages m_i^t , and MLP update function U_{mlp} to produce the updated intermediate embedding state $u_i^{t'}$:

$$\mathbf{m}_i^t = \mathbf{u}_i^t + \sum_{j=1}^S M_{\text{atn}}(\mathbf{u}_i^t, \mathbf{z}_j; \mathbf{W}_{\text{xy}}), \quad (10)$$

$$\mathbf{u}_i^{t'} = \mathbf{m}_i^t + U_{\text{mlp}}(\mathbf{m}_i^t; \mathbf{W}_{\text{xy}}). \quad (11)$$

The key advantage of Feature-to-Label message passing with attention is that each label node can attend differently on input elements (e.g. different words in an input sentence).

Updating Label Embeddings via Label-to-Label Message Passing

At this point, an independent prediction can be made for each label conditioned on \mathbf{x} using $\{\mathbf{u}_{1:L}^{t'}\}$. However, in order to consider label dependencies, we model interactions between the label nodes $\{\mathbf{u}_{1:L}^{t'}\}$ using Label-to-Label message passing and update them accordingly, as shown in the ‘‘Label-to-Label MP’’ block of Fig. 1. Given the exponentially large number of possible conditional dependencies, we use neural message passing as an efficient way to model such interactions, which has been shown to work well in practice for other tasks.

We assume there exist a label interaction graph $G_{\text{yy}} = (V_{\text{yy}}, E_{\text{yy}})$, $V_{\text{yy}} = \{y_{1:L}\}$, and E_{yy} includes all undirected pairwise edges connecting node y_i and node y_j . At this stage, we use another message passing module, MPNN_{yy} to pass messages between labels and update them. The label embedding $\mathbf{u}_i^{t'}$ is updated by a weighted combination through attention of all its neighbor label nodes $\{\mathbf{u}_{j \in \mathcal{N}(i)}^{t'}\}$.

MPNN_{yy} uses attention message function $M_{\text{atn}}^{t'}$ on all neighbor label embeddings $\{\mathbf{u}_{j \in \mathcal{N}(i)}^{t'}\}$ to produce message $\mathbf{m}_i^{t'}$, and MLP update function $U_{\text{mlp}}^{t'}$ to compute updated embedding $\mathbf{u}_i^{t'+1}$:

$$\mathbf{m}_i^{t'} = \mathbf{u}_i^{t'} + \sum_{j \in \mathcal{N}(i)} M_{\text{atn}}(\mathbf{u}_i^{t'}, \mathbf{u}_j^{t'}; \mathbf{W}_{\text{yy}}), \quad (12)$$

$$\mathbf{u}_i^{t'+1} = \mathbf{m}_i^{t'} + U_{\text{mlp}}(\mathbf{u}_i^{t'}, \mathbf{m}_i^{t'}; \mathbf{W}_{\text{yy}}). \quad (13)$$

If there exists a known label interaction graph G_{yy} , message $\mathbf{m}_i^{t'}$ for node i is computed using its neighboring nodes $j \in \mathcal{N}(i)$, where the neighbors $\mathcal{N}(i)$ are defined by the graph. If there is no known G_{yy} graph, we assume a fully connected graph, which means $\mathcal{N}(i) = \{j \in V_{\text{yy}}\}$ (including i).

Message Passing for Multiple Time Steps

To learn more complex relations among nodes, we compute a total of T time steps of updates. This is essentially a stack of T MPNN layers. In our implementation, the label embeddings are updated by MPNN_{xy} and MPNN_{yy} for T time steps to produce $\{\mathbf{u}_1^T, \mathbf{u}_2^T, \dots, \mathbf{u}_L^T\}$.

2.3 Readout Layer (MLC Predictions from the label embeddings)

After T updates to the label embeddings, the last module predicts each label $\{\hat{y}_1, \dots, \hat{y}_L\}$. A readout function R projects each of the L label embeddings \mathbf{u}_i^T using projection matrix $\mathbf{W}^o \in \mathbb{R}^{L \times d}$, where row $\mathbf{W}_i^o \in \mathbb{R}^d$ is the learned output vector for label i . The calculated vector of size $L \times 1$ is then fed through an element-wise sigmoid function to produce probabilities of each label being positive:

$$\hat{y}_i = R(\mathbf{u}_i^T; \mathbf{W}^o) = \text{sigmoid}(\mathbf{W}_i^o \mathbf{u}_i^T). \quad (14)$$

2.4 Model Details

Multi-head Attention. In order to allow a particular node to attend to multiple other nodes (or multiple groups of nodes) at once, LaMP uses multiple attention heads. Inspired by [52], we use K independent attention heads for each \mathbf{W} matrix during the message computation, where each matrix column $\mathbf{W}_j^{q,k}$ is of dimension d/K . The generated representations are concatenated (denoted by \parallel) and linearly transformed by matrix $\mathbf{W}^z \in \mathbb{R}^{d \times d}$. Multi-head attention changes message passing function M_{atn} , but update function U_{mlp} stays the same.

$$e_{ij}^{t,k} = (\mathbf{W}^{q,k} \mathbf{v}_i^t)^\top (\mathbf{W}^{u,k} \mathbf{v}_j^t) / \sqrt{d} \quad (15)$$

$$\alpha_{ij}^{t,k} = \frac{\exp(e_{ij}^{t,k})}{\sum_{j \in \mathcal{N}(i)} \exp(e_{ij}^{t,k})} \quad (16)$$

$$M_{\text{atn}}^k(\mathbf{v}_i^t, \mathbf{v}_j^t; \mathbf{W}) = \alpha_{ij}^{t,k} \mathbf{W}^{v,k} \mathbf{v}_j^t, \quad (17)$$

$$\mathbf{m}_i^t = \mathbf{v}_i^t + \left(\left\| \sum_{k=1}^K \left[\sum_{j \in \mathcal{N}(i)} M_{\text{atn}}^k(\mathbf{v}_i^t, \mathbf{v}_j^t; \mathbf{W}) \right] \right\| \right) \mathbf{W}^c \quad (18)$$

Matrices $\mathbf{W}^g, \mathbf{W}^u, \mathbf{W}^v, \mathbf{W}^r, \mathbf{W}^b, \mathbf{W}^c$, are not shared across time steps (but are shared across nodes).

Label Embedding Weight Sharing. To enforce each label’s input embedding to correspond to that particular label, the label embedding matrix weights \mathbf{W}^y are shared with the readout projection matrix \mathbf{W}^o . In other words, \mathbf{W}^y is used to produce the initial node vectors for G_{yy} , and then is used again to calculate the pre-sigmoid output values for each label, so $\mathbf{W}^o \equiv \mathbf{W}^y$. This was shown beneficial in Seq2Seq models for machine translation [38].

2.5 Loss Function

The final output of LaMP networks $\hat{\mathbf{y}}$ are trained using the mean binary cross entropy (BCE) over all outputs y_i . For one sample, given true binary label vector \mathbf{y} and predicted labels $\hat{\mathbf{y}}$, the output loss \mathcal{L}_{out} is:

$$\mathcal{L}_{\text{out}}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{L} \sum_{i=1}^L -(y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)) \quad (19)$$

The final outputs \hat{y}_i are computed from the final label node states \mathbf{u}_i^T (Eq. 14). However, since LaMP networks iteratively update the label nodes from $t = 0$ to T , we can “probe” the label nodes at each intermediate state from $t=1$ to $T-1$ and enforce an auxiliary loss on those states. To do this, we use the same matrix \mathbf{W}^o to extract the intermediate prediction \hat{y}_i^t at state t : $\hat{y}_i^t = R(\mathbf{u}_i^t; \mathbf{W}^o)$. We use the same BCE loss on these predictions to compute intermediate loss \mathcal{L}_{int} :

$$\mathcal{L}_{\text{int}}(\mathbf{y}, \hat{\mathbf{y}}^t) = \frac{1}{L} \sum_{i=1}^L -(y_i \log(\hat{y}_i^t) + (1 - y_i) \log(1 - \hat{y}_i^t)). \quad (20)$$

We note that the intermediate predictions \hat{y}_i^t are computed for both \mathbf{u}_i^t (after Label-to-Label message passing), as well as \mathbf{u}_i^t (after Feature-to-Label message passing). The final loss is a combination of both the original and intermediate, where the intermediate loss is weighted by λ :

$$\mathcal{L}_{\text{LaMP}} = \mathcal{L}_{\text{out}}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \sum_{t=1}^{T-1} \mathcal{L}_{\text{int}}(\mathbf{y}, \hat{\mathbf{y}}^t) \quad (21)$$

In LaMP networks, $p(y_i | \{y_{j \neq i}\}, \mathbf{z}_{1:S}; \mathbf{W})$ is approximated by jointly representing $\{y_{1:L}\}$ using message passing from $\{\mathbf{z}_{1:S}\}$ and from the embeddings of all neighboring labels $\{y_{j \in \mathcal{N}(i)}\}$.

2.6 LaMP Variation: Input Encoding with Feature Message Passing (FMP)

Thus far, we have assumed that we use the raw feature embeddings $\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_S\}$ to pass messages to the labels. However, we could also update the feature embeddings before they are passed to the label nodes by modelling the interactions between features.

For a particular input \mathbf{x} , we first assume that the input features $\{x_{1:S}\}$ are nodes on a graph, G_{xx} . $G_{\text{xx}} = (V_{\text{xx}}, E_{\text{xx}})$, $V_{\text{xx}} = \{x_{1:S}\}$, and E includes all undirected pairwise edges connecting node \mathbf{x}_i and node \mathbf{x}_j . MPNN_{xx} , parameterized by \mathbf{W}_{xx} , is used to pass messages between the input embeddings in order to update their states. Nodes on G_{xx} are represented as embedding vectors $\{\mathbf{z}_1^t, \mathbf{z}_2^t, \dots, \mathbf{z}_S^t\}$, where the initial states $\{\mathbf{z}_{1:S}^0\}$ are obtained using embedding matrix \mathbf{W}^x on input components $\{x_1, x_2, \dots, x_S\}$. The embeddings are then updated by MPNN_{xx} using message passing for T time steps to produce $\{\mathbf{z}_1^T, \mathbf{z}_2^T, \dots, \mathbf{z}_S^T\}$.

To update input embedding z_i^t , MPNN_{xx} uses attention message function M_{atn}^t (Eq. 6) on all neighboring input embeddings $\{z_j^t\}_{j \in \mathcal{N}(i)}$ to produce messages m_i^t , and MLP update function U_{mlp} (Eq. 9) to produce updated embedding z_i^{t+1} :

$$m_i^t = z_i^t + \sum_{j \in \mathcal{N}(i)} M_{\text{atn}}(z_i^t, z_j^t; \mathbf{W}_{\text{xx}}), \quad (22)$$

$$z_i^{t+1} = m_i^t + U_{\text{mlp}}(m_i^t; \mathbf{W}_{\text{xx}}). \quad (23)$$

If there exists a known G_{xx} graph, message m_i^t for node i is computed using its neighboring nodes $j \in \mathcal{N}(i)$, where the neighbors $\mathcal{N}(i)$ are defined by the graph. If there is no known graph, we assume a fully connected G_{xx} graph, which means $\mathcal{N}(i) = \{j \in V_{\text{xx}}\}$. Inputs with a sequential ordering can be modelled as a fully connected graph using positional embeddings [4].

In summary, MPNN_{xx} is used to update input feature nodes $\{z_{1:S}^t\}$ by passing messages within the feature-interaction graph. MPNN_{xy} is used to update output label nodes $\{u_{1:L}^t\}$ by passing messages from the features to labels (from input nodes $\{z_{1:S}^t\}$ to output nodes $\{u_{1:L}^t\}$). MPNN_{yy} is used to update output label nodes $\{u_{1:L}^t\}$ by passing messages within the label-interaction graph (between label nodes). Once messages have been passed to update the feature and label nodes for T integrative updates, a readout function R is then used on the label nodes to make a binary classification prediction on each label, $\{\hat{y}_1, \hat{y}_2, \dots, \hat{y}_L\}$. Figure 1 shows the LaMP network without the feature-interaction graph.

2.7 Advantages of LaMP Models

Efficiently Handling Dense Label Predictions. It is known that autoregressive models such as RNN Seq2Seq suffer from the propagation of errors over the sequential positive label predictions. This makes it difficult for these models to handle dense, or many positive label, samples. In addition, autoregressive models require a time consuming post-processing step such as beam search to obtain the optimal label set. Lastly, autoregressive models require a predefined label ordering for training the sequential prediction, which can lead to instabilities at testing time [54].

Motivated by the drawbacks of autoregressive models for MLC, the proposed LaMP model removes the reliance on sequential predictions, beam search, and a chosen label ordering, while still modelling the label dependencies. This is particularly beneficial when the number of positive output labels is large (i.e. dense). LaMP networks predict the output *set* of labels all at once, which is made possible by the fact that inference doesn't use a probabilistic chain, but there is still a representation of label dependencies via label to label attention. As an additional benefit, as noted by [5], it may be useful to maintain 'soft' predictions for each label in MLC. This is a major drawback of the PCC models which make 'hard' predictions of the positive labels, defaulting all other labels to 0.

Structure Agnostic. Many input or output types are instances where the relational structure is not made explicit, and must be inferred or assumed [4]. LaMP networks allow for greater flexibility of both input structures (known structure such as sequence or graph, or unknown such as tabular), as well as output structures (e.g., known graph vs unknown structure). To the best of our knowledge, this is the first work to use MPNNs to *infer* the relational structure of the data by using attention mechanisms.

Interpretability. Our formulation of LaMP allows us to visualize predictions in several different ways. First, since predictions are made in an iterative manner via graph update steps, we can "probe" each label's state at each step to get intermediate predictions. Second, we can visualize the attention weights which automatically learn the relational structure. Combining these two visualization methods allows us to see how the predictions change from the initial predictions given only the input sequence to the final state where messages have been passed from other labels, leading us to better insights for specific MLC samples.

2.8 Connecting to Related Topics

Structured Output Predictions. The use of graph attention in LaMP models is closely connected to the literature of structured output prediction for MLC. [14] used conditional random fields (CRFs)

	Reuters	Bibtex	Bookmarks	Delicious	RCV1	TFBS	SIDER	NUSWIDE
FastXML [37]	-	-	-	-	0.841	-	-	-
Madjarov [32]	-	0.434	0.257	0.343	-	-	-	-
SPEN [5]	-	0.422	0.344	0.375	-	-	-	-
RNN Seq2Seq [36]	0.894	0.393	0.362	0.320	0.890	0.249	0.356	0.329
Emb + MLP	0.854	0.363	0.368	0.371	0.865	0.167	0.766	0.371
Emb + LaMP _{el}	0.859	0.379	0.351	0.358	0.868	0.289	0.767	0.376
Emb + LaMP _{fc}	0.896	0.427	0.376	0.368	0.871	0.319	0.763	0.376
Emb + LaMP _{pr}	0.895	0.424	0.373	0.366	0.870	0.317	0.765	0.372
FMP + LaMP _{el}	0.883	0.435	0.375	0.369	0.887	0.310	0.766	-
FMP + LaMP _{fc}	0.906	0.445	0.389	0.372	0.889	0.321	0.764	-
FMP + LaMP _{pr}	0.902	0.447	0.386	0.372	0.887	0.321	0.766	-

Table 1: **ebF1 Scores** across all 8 datasets

[29] to model dependencies among labels and features for MLC by learning a distribution over pairs of labels to input features, but these are limited to pairwise dependencies.

To overcome the naive pairwise dependency constraint of CRFs, structured prediction energy networks (SPENs) [5] and related methods [50, 20] locally optimize an unconstrained structured output. In contrast to SPENs which use an iterative refinement of the output label predictions, our method is a simpler feed forward block to make predictions in one step, yet still models dependencies through attention mechanisms on embeddings, which gives the added interpretability benefit.

Multi-label Classification By Modeling Label Interaction Graphs. [19] formulate MLC using a label graph and they introduced a conditional dependency SVM where they first trained separate classifiers for each label given the input and all other true labels and used Gibbs sampling to find the optimal label set. The main drawback is that this method does not scale to a large number of labels. [42] proposes a method to label the pairwise edges of randomly generated label graphs, and requires some chosen aggregation method over all random graphs. The authors introduce the idea that variation in the graph structure shifts the inductive bias of the base learners. Our fully connected label graph with attention on the neighboring nodes can be regarded as a form of graph ensemble learning [22]. [11] use graph neural networks for MLC, but focus on graph inputs. They do not explicitly model label-to-label dependencies, thus resulting in a worse performance than LaMP.

Graph Neural Networks (GNNs). Passing embedding messages from node to neighbor nodes connects to a large body of literature on graph neural networks [4] and embedding models for structures [8].

The key idea is that instead of conducting probabilistic operations (e.g., product or re-normalization), the proposed models perform nonlinear function mappings in each step to learn feature representations of structured components. [15, 53, 3] all follow similar ideas to pass the embedding from node to neighbor nodes or neighbor edges.

There have been many recent works extending the basic GNN framework to update nodes using various message passing, update, and readout functions [26, 21, 31, 24, 59, 3, 15, 12]. We refer the readers to [4] for a survey. However, none of these have used GNNs for MLC. In addition, none of these have attempted to learn the graph structure by using neural attention on fully connected graphs.

3 Experiments

We validate our model on eight real world MLC datasets. These datasets vary in the number of samples, number of labels, input type (sequential, tabular, graph, vector), and output type (unknown, known label graph). They also cover a wide spectrum of input data types, including: raw English text (sequential form), binary word vector (tabular form), drug molecules (graph form), and images (vector form). Data statistics are in Table 6 and Section 8.1. Due to the space limit, we move the details of evaluation metrics to Section 8.2 and the hyper-parameters to Section 8.3. Details of previous results from the state-of-the-art baselines are in Section 8.4.

3.1 LaMP Variations

For LaMP models, we use two variations of input features, and three variations of Label-to-Label Message Passing. For input features, we use (1) **Emb**, which is the raw learned feature embeddings

	Reuters	Bibtex	Bookmarks	Delicious	RCV1	TFBS	SIDER	NUSWIDE
FastXML [37]	-	-	-	-	0.847	-	-	-
SVM [9]	0.787	-	-	-	-	-	-	-
GAML [11]	-	-	0.333	-	-	-	-	0.398
Madjarov [32]	-	0.462	0.268	0.339	-	-	-	-
RNN Seq2Seq [36]	0.858	0.384	0.329	0.329	0.884	0.311	0.389	0.418
Emb + MLP	0.835	0.389	0.349	0.385	0.855	0.218	0.795	0.465
Emb + LaMP _{el}	0.842	0.413	0.334	0.372	0.858	0.401	0.797	0.472
Emb + LaMP _{fc}	0.871	0.458	0.363	0.379	0.859	0.449	0.797	0.470
Emb + LaMP _{pr}	0.877	0.462	0.363	0.380	0.859	0.448	0.798	0.468
FMP + LaMP _{el}	0.870	0.455	0.355	0.381	0.877	0.445	0.795	-
FMP + LaMP _{fc}	0.886	0.465	0.373	0.384	0.877	0.450	0.795	-
FMP + LaMP _{pr}	0.889	0.473	0.371	0.386	0.877	0.449	0.797	-

Table 2: **miF1 Scores** across all 8 datasets

of dimension d , and (2) **FMP**² which is the updated state of each feature embedding after 2 layers of Feature Message Passing, as explained in 2.6. For each of the two input feature variations, we use three variations of the label graph which Label-to-Label Message Passing uses to update the labels given the input features, explained as follows.

LaMP_{el} uses an edgeless label graph and messages are not passed between labels, assuming no label dependencies.

LaMP_{fc} uses a fully connected label graph where each label is able to attend to all other labels (including itself) in order to compute the messages.

LaMP_{pr} uses a prior label graph where each label is able to attend to only other labels from the known label graph (including itself) in order to compute the messages. For RCV1, we use the known tree label structure, and for TFBS we use known protein-protein interactions (PPI) from [44]. For all other datasets, we create a graph where we place an edge on the adjacency matrix for all labels that co-occur in any sample for the training set. This is summarized in the last column of Appendix Table 5.

3.2 Performance Evaluation

ebF1. Table 1 shows the most commonly used evaluation, example-based F1 (ebF1) scores, for the seven datasets. LaMP outperforms the baseline MLP models which assume no label dependencies, as well as RNN Seq2Seq, which models label dependencies using a classifier chain. More importantly, we compare using an output graph with no edges (LaMP_{el}), which assumes no label dependencies vs. an output graph with edges (LaMP_{fc}). The two models have the same architecture and number of parameters, with the only thing varying being the message passing between label nodes. We can see that for most datasets, modelling label dependencies using LaMP_{fc} does in fact help. We found that using a known prior label structure (LaMP_{pr}) did not improve the results significantly. LaMP_{fc} predictions produced an average 1.8% ebF1 score increase over the independent LaMP_{el} predictions. LaMP_{pr} resulted in an average 1.7% ebF1 score increase over LaMP_{el}. When comparing to the MLP baseline, LaMP_{fc} and LaMP_{pr} produced an average 18.5% and 18.4% increase, respectively.

miF1. While high ebF1 scores indicate strong average F1 scores over all samples, the label-based Micro-averaged F1 (miF1) scores indicate strong results on the most frequent labels. Table 2 shows the miF1 scores, for the all datasets. LaMP_{fc} produced an average 1.6% miF1 score increase over the independent LaMP_{el}. LaMP_{pr} produced an average 1.8% miF1 score increase over LaMP_{el}. When comparing to the MLP baseline, LaMP_{fc} and LaMP_{pr} resulted in an average 20.2% and 20.5% increase, respectively.

maF1. Contrarily, high label-based Macro-averaged F1 (maF1) scores indicate strong results on less frequent labels. Table 2 shows maF1 scores, which show the strongest improvement of LaMP_{fc} and LaMP_{pr} variation over independent predictions. LaMP_{fc} resulted in an average 2.4% maF1 score increase over the independent LaMP_{el}. LaMP_{pr} produced an average 2.1% maF1 score increase over LaMP_{el}. This indicates that Label-to-Label message passing can help boost the accuracy of rare label predictions. When comparing to Emb + MLP, LaMP_{fc} and LaMP_{pr} produced an average 57.0% and 56.7% increase, respectively.

²For NUS-WIDE, since we use the 128-dimensional cVLAD features as input to compare to [11], we cannot use the **FMP** method.

	Reuters	Bibtex	Bookmarks	Delicious	RCV1	TFBS	SIDER	NUSWIDE
SVM [9]	0.468	-	-	-	-	-	-	-
FastXML [37]	-	-	-	-	0.592	-	-	-
GAML [11]	-	-	0.217	-	-	-	-	0.114
Madjarov [32]	-	0.316	0.119	0.142	-	-	-	-
RNN Seq2Seq [36]	0.457	0.282	0.237	0.166	0.741	0.210	0.207	0.143
Emb + MLP	0.366	0.275	0.248	0.180	0.667	0.094	0.665	0.173
Emb + LaMP _{el}	0.476	0.308	0.229	0.176	0.680	0.326	0.666	0.198
Emb + LaMP _{fc}	0.547	0.366	0.271	0.192	0.691	0.362	0.663	0.203
Emb + LaMP _{pr}	0.560	0.372	0.267	0.192	0.698	0.365	0.663	0.196
FMP + LaMP _{el}	0.508	0.353	0.266	0.192	0.742	0.368	0.664	-
FMP + LaMP _{fc}	0.520	0.371	0.286	0.195	0.743	0.364	0.668	-
FMP + LaMP _{pr}	0.517	0.376	0.280	0.196	0.740	0.364	0.664	-

Table 3: **maF1 Scores** across all 8 datasets

Other Metrics. Due to space constraints, we report subset accuracy in Appendix (supplementary) Table 7. RNN Seq2Seq models mostly perform all other metrics for this metric since they are trained to maximize it[36]. However, for all other metrics, RNN Seq2Seq does not perform as well, concluding that for most applications, PCC models aren’t necessary. We also report Hamming Accuracy in Appendix Table 8, and we note that LaMP networks outperform or perform similarly to baseline methods, but we observe that this metric is mostly unhelpful.

Metrics Performance Summary. While LaMP does not explicitly model label dependencies as autoregressive or structured prediction models do, the attention weights do learn some dependencies among labels (Section 3.3). This is indicated by the fact that LaMP, which uses Label-to-Label attention, mostly outperforms the ones which don’t, indicating that it is learning label dependencies.

Speed. LaMP results in a mean of 1.7x and 5.0x training and testing speedups, respectively, over the previous state-of-the-art probabilistic MLC method, RNN Seq2Seq. Speedups over RNN Seq2Seq model are shown in Table 3.3.

3.3 Interpretability Evaluation

The structure of LaMP networks allows for three different types of visualization methods to understand how the network predicts each label. We explain the three types here and show the results for a sample from the Bookmarks dataset using the FMP + LaMP_{fc} model.

Intermediate Output Prediction. One advantage of the multi step formulation of label embedding updates is that it gives us the ability to probe the state of each label at intermediate steps and view the model’s predictions at those steps. To do this, we use the readout function R on each intermediate label embeddings state u_i^t to find the probability that the label embedding would predict a positive label. In other words, this is the post-sigmoid output of the readout function of each embedding $R(u_i^t; \mathbf{W}^\circ)$ at each step $t = 1, \dots, T$. We note that each step contains two stages: $t.1$ is the output after the Feature-to-Label message passing, and $t.2$ is output after the Label-to-Label message passing. The output after the second stage of the final step (i.e. $T.2$) is the model’s final output.

Figure 2 (a.) shows the intermediate prediction outputs from the $T = 2$ step model. On the horizontal axis are a selected subset of all possible labels, with the red colored axis labels being all true positive labels. On the vertical axis, each row represents one of the label embedding states in the $T = 2$ step model. Each cell represents the readout function’s prediction for each label embedding’s state. The brighter the grid cell, the more likely that label is positive at the current stage. Starting from the bottom, the first row (1.1) shows the prediction of each label after the first Feature-to-Label message passing. The second row (1.2) shows the prediction of each label after the first Label-to-Label message passing. This is then repeated once more in (2.1) and (2.2) for the second layer’s output states, where the final output, 2.2 is the network’s final output predictions. The most important aspect of this figure is that we can see the labels “design”, “html”, and “web design”, all change from weakly positive to strongly positive after the first Label-to-Label message passing step (row 1.2). In other words, this indicates that these labels change to a strongly positive prediction by passing messages between each other.

Label-to-Feature Attention. While the iterative prediction visualization shows how the model updates its prediction of each label, it doesn’t explicitly show how or why. To understand why each label changes its predictions, we first look at the Feature-to-Label attention, which tells us the input nodes that each label node attends to in order to update its state (and thus producing the predictions

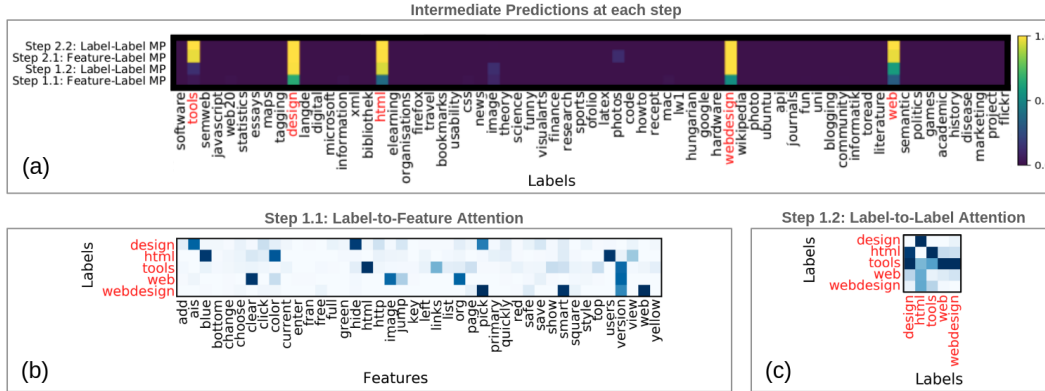


Figure 2: (a) **Visualization of Model Predictions and Attention Weights** Intermediate Predictions: this shows the readout function predictions for each intermediate state in the two update steps. (b) **Label-to-Feature Attention Weights** for the first step of Feature-to-Label message passing (1.1). (c) **Label-to-Label Attention Weights** for the first step of Label-to-Label message passing (1.2).

Dataset	Training	Testing
Reuters	0.788 (1.5x)	0.116 (2.1x)
Bibtex	0.376 (2.1x)	0.080 (2.1x)
Delicious	3.172 (1.1x)	0.473 (3.2x)
Bookmarks	9.664 (1.2x)	1.849 (1.3x)
RCV1	98.346 (1.2x)	1.003 (1.7x)
TFBS	187.14 (2.5x)	13.04 (4.2x)
NUS-WIDE	3.201 (1.2x)	0.921 (8.0x)
SIDER	0.027 (2.5x)	0.003 (21x)

Table 4: **Speed.** Each column shows training or testing speed for LaMP in minutes per epoch. Speedups over RNN Seq2Seq are in parentheses. Since LaMP does not depend on sequential prediction, we see a drastic speedup, especially during testing where RNN Seq2Seq requires beam search.

in Figure 2 (a.)). Figure 2 (b.) shows us which input nodes (i.e. features) each of the positive label attends to in order to make its first update step 1.1. The colors represent the post-softmax attention weight (summed over the 4 attention heads), with the darker cells representing high attention. In this example, we can see that the “web design” label attends to the “pick”, “smart”, and “version” features, but as we can see from the first row of Figure 2 (a.), prediction for the current state of the “web design” label isn’t very strong yet.

Label-to-Label Attention. Label-to-Feature attention shows us the input nodes that each label node attends to in order to make its first update, but the second step of the label graph update is the Label-to-Label message passing step where labels are updated according to the states of all other nodes after the first Feature-to-Label message passing. Figure 2 (c.) shows us the first Label-to-Label attention stage 1.2 where each label node can attend to the other label nodes in order to update its state. Here we show only the Label-to-Label attention for the positive labels in this example. We then look at the second row of Figure 2 (a.) which shows the model’s prediction of each label node after the Label-to-Label attention. The interesting thing to note here is we can see many of the true positive labels change their state to positive after the positive labels attend to each other during the Label-to-Label attention step, indicating that dependencies are learned.

Attention weights for the second step $t = 2$ are not as interpretable since they model higher order interactions. We have added these plots in Appendix Fig. 3.

4 Conclusion

In this work we present Label Message Passing (LaMP) Networks which achieve better than, or close to the same accuracy as previous methods across five metrics and seven datasets. In addition, the iterative label embedding updates with attention of LaMP provide a straightforward way to shed light on the model’s predictions and allow us to extract three forms of visualizations, including conditional label dependencies which influence MLC classifications.

References

[1] Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)

- [2] Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014)
- [3] Battaglia, P., Pascanu, R., Lai, M., Rezende, D.J., et al.: Interaction networks for learning about objects, relations and physics. In: Advances in neural information processing systems. pp. 4502–4510 (2016)
- [4] Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al.: Relational inductive biases, deep learning, and graph networks. arXiv:1806.01261 (2018)
- [5] Belanger, D., McCallum, A.: Structured prediction energy networks. In: International Conference on Machine Learning. pp. 983–992 (2016)
- [6] Bhatia, K., Jain, H., Kar, P., Varma, M., Jain, P.: Sparse local embeddings for extreme multi-label classification. In: Neural Information Processing Systems. pp. 730–738 (2015)
- [7] Consortium, E.P., et al.: An integrated encyclopedia of dna elements in the human genome. *Nature* **489**(7414), 57 (2012)
- [8] Dai, H., Dai, B., Song, L.: Discriminative embeddings of latent variable models for structured data. In: International Conference on Machine Learning. pp. 2702–2711 (2016)
- [9] Debole, F., Sebastiani, F.: An analysis of the relative hardness of reuters-21578. *American Society for Information Science and Technology* **56**(6), 584–596 (2005)
- [10] Dembczynski, K., Cheng, W., Hüllermeier, E.: Bayes optimal multilabel classification via probabilistic classifier chains. (2010)
- [11] Do, K., Tran, T., Nguyen, T., Venkatesh, S.: Attentional multilabel learning over graphs—a message passing approach. arXiv preprint arXiv:1804.00293 (2018)
- [12] Duvenaud, D.K., Maclaurin, D., Iparraguirre, J., Bombarell, R., Hirzel, T., Aspuru-Guzik, A., Adams, R.P.: Convolutional networks on graphs for learning molecular fingerprints. In: Advances in neural information processing systems. pp. 2224–2232 (2015)
- [13] Elisseeff, A., Weston, J.: A kernel method for multi-labelled classification. In: Advances in neural information processing systems. pp. 681–687 (2002)
- [14] Ghamrawi, N., McCallum, A.: Collective multi-label classification. In: 14th ACM international conference on Information and knowledge management. pp. 195–200. ACM (2005)
- [15] Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. arXiv preprint arXiv:1704.01212 (2017)
- [16] Godbole, S., Sarawagi, S.: Discriminative methods for multi-labeled classification. In: Pacific-Asia conference on knowledge discovery and data mining. pp. 22–30. Springer (2004)
- [17] Graves, A.: Generating sequences with recurrent neural networks. arXiv preprint arXiv:1308.0850 (2013)
- [18] Gu, J., Bradbury, J., Xiong, C., Li, V.O., Socher, R.: Non-autoregressive neural machine translation. arXiv preprint arXiv:1711.02281 (2017)
- [19] Guo, Y., Gu, S.: Multi-label classification using conditional dependency networks. In: IJCAI Proceedings-International Joint Conference on Artificial Intelligence. vol. 22, p. 1300 (2011)
- [20] Gygli, M., Norouzi, M., Angelova, A.: Deep value networks learn to evaluate and iteratively refine structured outputs. arXiv preprint arXiv:1703.04363 (2017)
- [21] Hamilton, W.L., Ying, R., Leskovec, J.: Representation learning on graphs: Methods and applications. arXiv preprint arXiv:1709.05584 (2017)
- [22] Hara, K., Saitoh, D., Shouno, H.: Analysis of dropout learning regarded as ensemble learning. In: International Conference on Artificial Neural Networks. pp. 72–79. Springer (2016)
- [23] Katakis, I., Tsoumakas, G., Vlahavas, I.: Multilabel text classification for automated tag suggestion
- [24] Kearnes, S., McCloskey, K., Berndl, M., Pande, V., Riley, P.: Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design* **30**(8), 595–608 (2016)
- [25] Kingma, D., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)

- [26] Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
- [27] Koehn, P., Knowles, R.: Six challenges for neural machine translation. arXiv preprint arXiv:1706.03872 (2017)
- [28] Kuhn, M., Letunic, I., Jensen, L.J., Bork, P.: The sider database of drugs and side effects. *Nucleic acids research* **44**(D1), D1075–D1079 (2015)
- [29] Lafferty, J., McCallum, A., Pereira, F.C.: Conditional random fields: Probabilistic models for segmenting and labeling sequence data (2001)
- [30] Lewis, D.D., Yang, Y., Rose, T.G., Li, F.: Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research* **5**(Apr), 361–397 (2004)
- [31] Li, Y., Tarlow, D., Brockschmidt, M., Zemel, R.: Gated graph sequence neural networks. arXiv preprint arXiv:1511.05493 (2015)
- [32] Madjarov, G., Kocev, D., Gjorgjevikj, D., Džeroski, S.: An extensive experimental comparison of methods for multi-label learning. *Pattern recognition* **45**(9), 3084–3104 (2012)
- [33] McCallum, A.: Multi-label text classification with a mixture model trained by em
- [34] Montañes, E., Senge, R., Barranquero, J., Quevedo, J.R., del Coz, J.J., Hüllermeier, E.: Dependent binary relevance models for multi-label classification. *Pattern Recognition* **47**(3), 1494–1508 (2014)
- [35] Nam, J., Kim, J., Mencía, E.L., Gurevych, I., Fürnkranz, J.: Large-scale multi-label text classification revisiting neural networks. In: *Joint european conference on machine learning and knowledge discovery in databases*. pp. 437–452. Springer (2014)
- [36] Nam, J., Mencía, E.L., Kim, H.J., Fürnkranz, J.: Maximizing subset accuracy with recurrent neural networks in multi-label classification. In: *Advances in Neural Information Processing Systems*. pp. 5419–5429 (2017)
- [37] Prabhu, Y., Varma, M.: Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. pp. 263–272. ACM (2014)
- [38] Press, O., Wolf, L.: Using the output embedding to improve language models. arXiv preprint arXiv:1608.05859 (2016)
- [39] Read, J., Pfahringer, B., Holmes, G., Frank, E.: Classifier chains for multi-label classification. *Machine Learning and Knowledge Discovery in Databases* pp. 254–269 (2009)
- [40] Read, J., Pfahringer, B., Holmes, G., Frank, E.: Classifier chains for multi-label classification. *Machine learning* **85**(3), 333 (2011)
- [41] Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. *IEEE Transactions on Neural Networks* **20**(1), 61–80 (2009)
- [42] Su, H., Rousu, J.: Multilabel classification through random graph ensembles. In: *Asian Conference on Machine Learning*. pp. 404–418 (2013)
- [43] Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: *Advances in neural information processing systems*. pp. 3104–3112 (2014)
- [44] Szklarczyk, D., Morris, J.H., Cook, H., Kuhn, M., Wyder, S., Simonovic, M., Santos, A., et al.: The string database in 2017: quality-controlled protein–protein association networks, made broadly accessible. *Nucleic acids research* p. gkw937 (2016)
- [45] Tsochantaridis, I., Joachims, T., Hofmann, T., Altun, Y.: Large margin methods for structured and interdependent output variables. *JMLR* **6**(Sep), 1453–1484 (2005)
- [46] Tsoumakas, G., Katakis, I.: Multi-label classification: An overview. *International Journal of Data Warehousing and Mining* **3**(3) (2006)
- [47] Tsoumakas, G., Katakis, I., Vlahavas, I.: Effective and efficient multilabel classification in domains with large number of labels
- [48] Tsoumakas, G., Katakis, I., Vlahavas, I.: Mining multi-label data. In: *Data mining and knowledge discovery handbook*, pp. 667–685. Springer (2009)

- [49] Tsoumakas, G., Vlahavas, I.: Random k-labelsets: An ensemble method for multilabel classification. In: European conference on machine learning. pp. 406–417. Springer (2007)
- [50] Tu, L., Gimpel, K.: Learning approximate inference networks for structured prediction. arXiv preprint arXiv:1803.03376 (2018)
- [51] Ueda, N., Saito, K.: Parametric mixture models for multi-labeled text. In: Advances in neural information processing systems. pp. 737–744 (2003)
- [52] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. pp. 6000–6010 (2017)
- [53] Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. arXiv preprint arXiv:1710.10903 (2017)
- [54] Vinyals, O., Bengio, S., Kudlur, M.: Order matters: Sequence to sequence for sets. arXiv preprint arXiv:1511.06391 (2015)
- [55] Wang, J., Yang, Y., Mao, J., Huang, Z., Huang, C., Xu, W.: Cnn-rnn: A unified framework for multi-label image classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2285–2294 (2016)
- [56] Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: International Conference on Machine Learning. pp. 2048–2057. <http://www.jmlr.org/proceedings/papers/v37/xuc15.pdf>
- [57] Yeh, C.K., Wu, W.C., Ko, W.J., Wang, Y.C.F.: Learning deep latent space for multi-label classification. In: AAAI. pp. 2838–2844 (2017)
- [58] Zhang, M.L., Zhou, Z.H.: A k-nearest neighbor based algorithm for multi-label classification. In: Granular Computing, 2005 IEEE International Conference on. vol. 2, pp. 718–721. IEEE (2005)
- [59] Zheng, D., Luo, V., Wu, J., Tenenbaum, J.B.: Unsupervised learning of latent physical properties using perception-prediction networks. arXiv preprint arXiv:1807.09244 (2018)

5 Appendix: MLC Background

5.1 Background of Multi-Label Classification:

MLC has a rich history in text [33, 51], images [46, 13], bioinformatics [46, 13], and many other domains. MLC methods can roughly be broken into several groups, which are explained as follows.

Label powerset models (LP) [49, 40], classify each input into one label combination from the set of all possible combinations $\mathcal{Y} = \{\{1\}, \{2\}, \dots, \{1, 2, \dots, L\}\}$. LP explicitly models the joint distribution by predicting the one subset of all positive labels. Since the label set Y grows exponentially in the number of total labels (2^L), classifying each possible label set is intractable for a modest L . In addition, even in small L tasks, LP suffers from the “subset scarcity problem” where only a small amount of the label subsets are seen during training, leading to bad generalization.

Binary relevance (BR) methods predict each label separately as a logistic regression classifier for each label [58, 16]. The naïve approach to BR prediction is to predict all labels independently of one another, assuming no dependencies among labels. That is, BR uses the following conditional probability parameterized by learned weights W :

$$P_{BR}(Y|X; W) = \prod_{i=1}^L p(Y_i|X_{1:S}; W) \tag{24}$$

Probabilistic classifier chain (PCC) methods [10, 39] are autoregressive models that estimate the true joint probability of output labels given the input by using the chain rule, predicting one label at a time:

$$P_{PCC}(Y|X; W) = \prod_{i=1}^L p(Y_i|Y_{1:i-1}, X_{1:S}; W) \tag{25}$$

Two issues with PCC models are that inference is very slow if L is large, and the errors propagate as L increases [34]. To mitigate the problems with both LP and PCC methods, one solution is to only predict the true labels in the LP subset. In other words, only predicting the positive labels (total of ρ for a particular sample) and ignoring all other labels, which we call PCC+. Similar to PCC, the joint probability of PCC+ can be computed as product of conditional probabilities, but unlike PCC, only $\rho < L$ terms are predicted as positive:

$$P_{PCC+}(Y|X; W) = \prod_{i=1}^{\rho} p(Y_{p_i}|Y_{p_{1:i-1}}, x_{1:S}; W) \tag{26}$$

This can be beneficial when the number of possible labels L is large, reducing the total number of prediction steps. However, in both PCC and PCC+, inference is done using beam search, which is a costly dynamic programming step to find the optimal prediction.

Recently, Recurrent neural network (RNN) based encoder-decoder models following PCC and PCC+ have shown state-of-the-art performance for solving MLC. However, the sequential nature of modeling label dependencies through an RNN limits its ability in parallel computation, predicting dense labels, and providing interpretable results.

The main drawback of classifier chain models is that their inherently sequential nature precludes parallelization during training and inference. This can be detrimental when there are a large number of positive labels as the classifier chain has to sequentially predict each label, and often requires beam search to obtain the optimal set. Aside from time-cost disadvantages, PCC methods have several other drawbacks. First, PCC methods require a defined ordering of labels for the sequential prediction, but MLC output labels are an unordered set, and the chosen order can lead to prediction instability [36]. Secondly, even if the optimal ordering is known, PCC methods struggle to accurately capture long-range dependencies among labels in cases where the number of positive labels is large (i.e., dense labels). For example, the Delicious dataset we used in the experiment has a median of 19 positive labels per sample, so it can be difficult to correctly predict the labels at the end of the prediction chain. Lastly, many real-world applications prefer interpretable predictors. For instance, in the task of predicting which proteins (labels) will bind to a DNA sequence based binding site, users

care about how a prediction is made and how the interactions among labels (proteins) influence the binding predictions. An important task is modelling what is known as “co-binding” effects, where one protein will *only* bind if there is another specific protein already binding, or similarly will not bind if there is another already binding.

LaMP methods approximate the following factored formulation, where $\mathcal{N}(Y_i)$ denotes the neighboring nodes of Y_i .

$$P_{G2G}(Y|X; W) = \prod_{i=1}^L p(Y_i | \{Y_{\mathcal{N}(Y_i)}\}, X_{1:S}; W). \quad (27)$$

5.2 Seq2Seq Models

In machine translation (MT), sequence-to-sequence (Seq2Seq) models have proven to be the superior method, where an encoder RNN reads the source language sentence into an encoder hidden state, and a decoder RNN translates the hidden state into a target sentence, predicting each word autoregressively [43]. [2] improved this model by introducing “neural attention” which allows the decoder RNN to “attend” to every encoder word at each step of the autoregressive translation.

Recently, [36] showed that, across several metrics, state-of-the-art MLC results could be achieved by using a recurrent neural network (RNN) based encoder-to-decoder framework for Equation 26 (PCC+). They use a Seq2Seq RNN model (Seq2Seq Autoregressive) which uses one RNN to encode x , and a second RNN to predict each positive label sequentially, until it predicts a ‘stop’ signal. This type of model seeks to maximize the ‘subset accuracy’, or correctly predict every label as its exact 0/1 value.

[52] eliminated the need for the recurrent network in MT by introducing the Transformer. Instead of using an RNN to model dependencies, the Transformer explicitly models pairwise dependencies among all of the features by using attention [2, 56] between signals. This speeds up training time because RNNs can’t be fully parallelized but, the transformer still uses an autoregressive decoder.

5.3 Drawbacks of Autoregressive Models

Seq2Seq MLC [36] uses an encoder RNN encoding elements of an input sequence, a decoder RNN predicting output labels one after another, and beam search that computes the probability of the next T predictions of labels and then chooses the solution with the max combined probability.

Autoregressive models have been proven effective for machine translation and MLC [43, 2, 36]. However, predictions must be made sequentially, eliminating parallelization. Also, beam search is typically used at test time to find optimal predictions. But beam search is limited by the time cost of large beams sizes, making it difficult to optimally predict many output labels [27].

In addition to speed constraints, beam search for autoregressive inference introduces a second drawback: initial wrong predictions will propagate when using a modest beam size (e.g. most models use a beam size of 5). This can lead to significant decreases in performance when the number of positive labels is large. For example, the Delicious dataset has a median of 19 positive labels per sample, and it can be very difficult to correctly predict the labels at the end of the prediction chain.

Autoregressive models are well suited for machine translation because these models mimic the sequential decoding process of real translation. However, for MLC, the output labels have no intrinsic ordering. While the joint probability of the output labels is independent of the label ordering via autoregressive based inference, the chosen ordering can make a difference in practice [54, 36]. Some ordering of labels must be used during training, and this chosen ordering can lead to unstable predictions at test time.

Our LaMP connects to [18] who removed the autoregressive decoder in MT with the Non-Autoregressive Transformer. In this model, the encoder makes a proxy prediction, called “fertilities”, which are used by the decoder to predict all translated words at once. The difference between their model and ours is that we have a constant label at each position, so we don’t need to marginalize over all possible labels at each position.

6 Appendix: Dataset Details

Dataset	Input Type	Domain	#Train	#Val	#Test	Labels (L)	Features	Prior Graph Structure
Reuters-21578	Sequential	Text	6,993	777	3,019	90	23,662	Co-occur
RCV1-V2	Sequential	Text	703,135	78,126	23,149	103	368,998	Tree
TFBS	Sequential	Biology	1,671,873	301,823	323,796	179	4	PPI
BibTex	Binary Vector	Text	4,377	487	2,515	159	1,836	Co-occur
Delicious	Binary Vector	Text	11,597	1,289	3,185	983	500	Co-occur
Bookmarks	Binary Vector	Text	48,000	12,000	27,856	208	368,998	Co-occur
NUS-WIDE	Vector	Image	129,431	32,358	107,859	85	128	Co-occur
SIDER	Graph	Drug	1,141	143	143	27	37	Co-occur

Table 5: **Dataset Statistics.** We use 7 well studied MLC datasets, plus our own TFBS protein binding dataset. Each dataset varies in the input type, number of samples, number of labels, and number of input features. The last column shows the prior graph structure type we explore for the LaMP_{pr} model.

Dataset	Mean Labels /Sample	Median Labels /Sample	Max Labels /Sample	Mean Samples /Label	Median Samples /Label	Max Samples /Label
Reuters-21578	1.23	1	15	106.50	18	2,877
RCV1-V2	3.21	3	17	24,362	7,250	363,991
TFBS	7.62	2	178	84,047	45,389	466,876
BibTex	2.38	2	28	72.79	54	689
Delicious	19.06	20	25	250.15	85	5,189
Bookmarks	2.03	1	44	584.67	381	4,642
NUS-WIDE	1.86	1	12	3721.7	1104	44255
SIDER	15.3	16	26	731.07	851	1185

Table 6: **Additional Dataset Statistics** Here we show additional statistics of datasets with respect to the specific number of labels for each dataset. This shows how each dataset has a varying degree of MLC difficulty regarding the number of labels which need to be predicted.

7 Appendix: Extra Metrics

Here we provide the results from an extra two metrics: subset accuracy and hamming accuracy.

	Reuters	Bibtex	Bookmarks	Delicious	RCV1	TFBS	SIDER	NUSWIDE
Madjarov	-	0.202	0.209	0.018	-	-	-	-
RNN Seq2Seq	0.837	0.195	0.273	0.016	0.6798	0.114	0.000	0.252
Emb + MLP	0.774	0.151	0.234	0.180	0.620	0.040	0.014	0.263
Emb + LaMP _{el}	0.757	0.141	0.214	0.176	0.619	0.077	0.014	0.268
Emb + LaMP _{fc}	0.813	0.171	0.234	0.192	0.630	0.086	0.007	0.269
Emb + LaMP _{pr}	0.813	0.169	0.232	0.192	0.621	0.087	0.007	0.267
FMP + LaMP _{el}	0.808	0.158	0.231	0.192	0.656	0.084	0.007	-
FMP + LaMP _{fc}	0.835	0.182	0.242	0.195	0.660	0.090	0.014	-
FMP + LaMP _{pr}	0.828	0.185	0.241	0.196	0.659	0.090	0.007	-

Table 7: **Subset Accuracy Scores** across all 7 datasets

	Reuters	Bibtex	Bookmarks	Delicious	RCV1	TFBS	SIDER	NUSWIDE
Madjarov	-	0.988	0.991	0.982	-	-	-	-
RNN Seq2Seq	0.996	0.985	0.990	0.980	0.9925	0.961	0.593	0.980
Emb + MLP	0.996	0.987	0.991	0.982	0.992	0.959	0.752	0.980
Emb + LaMP _{el}	0.996	0.987	0.991	0.982	0.992	0.963	0.750	0.980
Emb + LaMP _{fc}	0.997	0.988	0.992	0.982	0.992	0.964	0.752	0.980
Emb + LaMP _{pr}	0.997	0.988	0.991	0.982	0.992	0.964	0.751	0.980
FMP + LaMP _{el}	0.997	0.987	0.991	0.982	0.993	0.964	0.748	-
FMP + LaMP _{fc}	0.997	0.988	0.992	0.982	0.993	0.964	0.749	-
FMP + LaMP _{pr}	0.997	0.988	0.992	0.982	0.993	0.964	0.747	-

Table 8: **Hamming Accuracy** across all 7 datasets

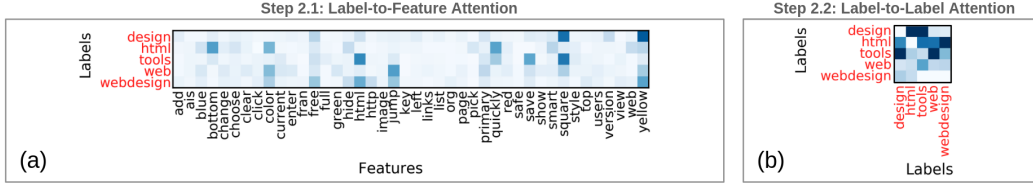


Figure 3: This shows the step $t = 2$ visualization results from Fig 2 (a). **2.1 Label-to-Feature Attention Weights** (b). **2.2 Label-to-Label Attention Weights** (c)

8 Appendix: More about Experiments

8.1 Datasets

We test our method against baseline methods on seven different multi-label sequence classification datasets. The datasets are summarized in Table 6. We use Reuters-21578 [30], Bibtex [48], Delicious [47], Bookmarks [23], RCV1-V2 [30], our own DNA protein binding dataset (TFBS) from [7], and SIDER [28], which is side effects of drug molecules. As shown in the table, each dataset has a varying number of samples, number of labels, positive labels per sample, and samples per label. For BibTex and Delicious, we use 10% of the provided training set for validation. For the TFBS dataset, we use 1 layer of convolution at the first layer to extract “words” from the DNA characters (A,C,G,T), as commonly done in deep learning models for DNA.

For datasets which have sequential ordering of the input components (Reuters, RCV1), we add a positional encoding to the word embedding as used in [52] (sine and cosine functions of different frequencies) to encode the location of each word in the sentence. For datasets with no ordering or graph structure (Bibtex, Delicious, Bookmarks, which use bag-of-word input representations) we do not use positional encodings. For inputs with an explicit graph representation (SIDER), we use the known graph structure.

8.2 Evaluation Metrics

Multi-label classification methods can be evaluated with many different metrics which each evaluate different strengths or weaknesses. We use the same 5 evaluation metrics from [36].

All of our autoregressive models predict only the positive labels before outputting a stop signal. This is a special case of PCC models, which have been shown to outperform the binary prediction of each label in terms of performance and speed. These models use beam search at inference time with a beam size of 5. For the non-autoregressive models, to convert the labels to $\{0, 1\}$ we chose the best threshold on the validation set from the same set of thresholds used in [50].

Example-based measures are defined by comparing the target vector \mathbf{y} to the prediction vector $\hat{\mathbf{y}}$. Subset Accuracy (ACC) requires an exact match of the predicted labels and the true labels: $\text{ACC}(\mathbf{y}, \hat{\mathbf{y}}) = \mathbb{I}[\mathbf{y} = \hat{\mathbf{y}}]$. Hamming Accuracy (HA) evaluates how many labels are correctly predicted in $\hat{\mathbf{y}}$: $\text{HA}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{L} \sum_{j=1}^L \mathbb{I}[y_j = \hat{y}_j]$. Example-based F1 (ebF1) measures the ratio of correctly predicted labels to the sum of the total true and predicted labels: $\frac{2 \sum_{j=1}^L y_j \hat{y}_j}{\sum_{j=1}^L y_j + \sum_{j=1}^L \hat{y}_j}$.

Label-based measures treat each label y_j as a separate two-class prediction problem, and compute the number of true positives (tp_j), false positives (fp_j), and false negatives (fn_j) for a label. Macro-averaged F1 (maF1) measures the label-based F1 averaged over all labels: $\frac{1}{L} \sum_{j=1}^L \frac{2tp_j}{2tp_j + fp_j + fn_j}$. Micro-averaged F1 (miF1) measures the label-based F1 averaged over each sample: $\frac{\sum_{j=1}^L 2tp_j}{\sum_{j=1}^L 2tp_j + fp_j + fn_j}$. High maF1 scores usually indicate high performance on less frequent labels. High miF1 scores usually indicate high performance on more frequent labels.

8.3 Model Hyperparameter Tuning

For all 7 datasets (Table 6), we use the same LaMP model with $T=2$ time steps, $d = 512$, and $K=4$ attention heads. We trained our models on an NVIDIA TITAN X Pascal with a batch size of 32.

We used Adam [25] with $\text{betas} = (0.9, 0.999)$, $\text{eps} = 1\text{e-}08$, and a learning rate of 0.0002 for each dataset. We used dropout of $p = 0.2$ for the smaller datasets (Reuters, Bibtext, SIDER), and dropout of $p = 0.1$ for all other datasets. The LaMP models also use layer normalization [1] around each of the attention and feedforward layers. All LaMP models are trained with the LaMP loss (Eq. 21). The hyperparameter λ is selected from the best performing value in $\{0, 0.1, 0.2, 0.3\}$ for each model. MLP models are trained with regular binary cross entropy (Eq. 19), and RNN Seq2Seq model are trained with cross entropy across all possible labels at each position. To convert the soft predictions into $\{0, 1\}$ values, we use the same thresholds in [5] and select the best one for each metric on the validation set. For the TFBS dataset, which uses DNA input sequences, we use one layer of convolution to get 512 dimensional embeddings as commonly done for deep neural network prediction tasks on DNA sequences.

8.4 Baseline Comparisons

Briefly, we compare against the following methods for all reported datasets and metrics. For those results named as “Madjarov”: we take the best method for each reported metric from [32] who compared 12 different types of models including SVMs, decision trees, boosting, classification rules, and neural networks. For results of “SPEN”: Structured prediction energy networks from [5]. For results of “SVM”: SVM method from the Reuters dataset authors [9]. For results of “FastXML”: Fast random forest model [37]. For results of “GAML”: graph attention for MLC from [11]. For “RNN Seq2Seq”: RNN Sequence to Sequence model from [35] which is a PCC model that predicts only the positive labels. For “Emb + MLP”: we use the mean embeddings of all input features as the input to a 4 layer multi-layer perceptron (MLP). This is a BR baseline which predicts all labels independently.