

Semi-Supervised Sequence Modeling with Cross-View Training

Kevin Clark¹ Minh-Thang Luong² Christopher D. Manning¹ Quoc V. Le²

¹Computer Science Department, Stanford University ²Google Brain

kevclark@cs.stanford.edu, thangluong@google.com, manning@cs.stanford.edu, qvl@google.com

Abstract

Unsupervised representation learning algorithms such as word2vec and ELMo improve the accuracy of many supervised NLP models, mainly because they can take advantage of large amounts of unlabeled text. However, the supervised models only learn from task-specific labeled data during the main training phase. We therefore propose Cross-View Training (CVT), a semi-supervised learning algorithm that improves the representations of a Bi-LSTM sentence encoder using a mix of labeled and unlabeled data. On labeled examples, standard supervised learning is used. On unlabeled examples, CVT teaches auxiliary prediction modules that see restricted views of the input (e.g., only part of a sentence) to match the predictions of the full model seeing the whole input. Since the auxiliary modules and the full model share intermediate representations, this in turn improves the full model. Moreover, we show that CVT is particularly effective when combined with multi-task learning. We evaluate CVT on five sequence tagging tasks, machine translation, and dependency parsing, achieving state-of-the-art results.¹

1 Introduction

Deep learning models work best when trained on large amounts of labeled data. However, acquiring labels is costly, motivating the need for effective semi-supervised learning techniques that leverage unlabeled examples. A widely successful semi-supervised learning strategy for neural NLP is pre-training word vectors (Mikolov et al., 2013). More recent work trains a Bi-LSTM sentence encoder to do language modeling and then incorporates its context-sensitive representations into supervised models (Dai and Le, 2015; Peters et al.,

2018). Such pre-training methods perform unsupervised representation learning on a large corpus of unlabeled data followed by supervised training.

A key disadvantage of pre-training is that the first representation learning phase does not take advantage of labeled data – the model attempts to learn generally effective representations rather than ones that are targeted towards a particular task. Older semi-supervised learning algorithms like self-training do not suffer from this problem because they continually learn about a task on a mix of labeled and unlabeled data. Self-training has historically been effective for NLP (Yarowsky, 1995; McClosky et al., 2006), but is less commonly used with neural models. This paper presents Cross-View Training (CVT), a new self-training algorithm that works well for neural sequence models.

In self-training, the model learns as normal on labeled examples. On unlabeled examples, the model acts as both a teacher that makes predictions about the examples and a student that is trained on those predictions. Although this process has shown value for some tasks, it is somewhat tautological: the model already produces the predictions it is being trained on. Recent research on computer vision addresses this by adding noise to the student’s input, training the model so it is robust to input perturbations (Sajjadi et al., 2016; Wei et al., 2018). However, applying noise is difficult for discrete inputs like text.

As a solution, we take inspiration from multi-view learning (Blum and Mitchell, 1998; Xu et al., 2013) and train the model to produce consistent predictions across different *views* of the input. Instead of only training the full model as a student, CVT adds auxiliary prediction modules – neural networks that transform vector representations into predictions – to the model and also trains them as students. The input to each student prediction module is a subset of the model’s intermediate rep-

¹Code will be made available at https://github.com/tensorflow/models/tree/master/research/cvt_text

representations corresponding to a restricted view of the input example. For example, one auxiliary prediction module for sequence tagging is attached to only the “forward” LSTM in the model’s first Bi-LSTM layer, so it makes predictions without seeing any tokens to the right of the current one.

CVT works by improving the model’s representation learning. The auxiliary prediction modules can learn from the full model’s predictions because the full model has a better, unrestricted view of the input. As the auxiliary modules learn to make accurate predictions despite their restricted views of the input, they improve the quality of the representations they are built on top of. This in turn improves the full model, which uses the same shared representations. In short, our method combines the idea of representation learning on unlabeled data with classic self-training.

CVT can be applied to a variety of tasks and neural architectures, but we focus on sequence modeling tasks where the prediction modules are attached to a shared Bi-LSTM encoder. We propose auxiliary prediction modules that work well for sequence taggers, graph-based dependency parsers, and sequence-to-sequence models. We evaluate our approach on English dependency parsing, combinatory categorial grammar supertagging, named entity recognition, part-of-speech tagging, and text chunking, as well as English to Vietnamese machine translation. CVT improves over previously published results on all these tasks. Furthermore, CVT can easily and effectively be combined with multi-task learning: we just add additional prediction modules for the different tasks on top of the shared Bi-LSTM encoder. Training a unified model to jointly perform all of the tasks except machine translation improves results (outperforming a multi-task ELMo model) while decreasing the total training time.


2 Cross-View Training

We first present Cross-View Training and describe how it can be combined effectively with multi-task learning. See Figure 1 for an overview of the training method.


2.1 Method

Let $\mathcal{D}_l = \{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$ represent a labeled dataset and $\mathcal{D}_{ul} = \{x_1, x_2, \dots, x_M\}$ represent an unlabeled dataset. We use $p_\theta(y|x_i)$ to denote the output distribution over classes pro-

Learning on a Labeled Example



Learning on an Unlabeled Example



Inputs Seen by Auxiliary Prediction Modules

Auxiliary 1: They traveled to _____
 Auxiliary 2: They traveled to **Washington** _____
 Auxiliary 3: _____ **Washington** by plane
 Auxiliary 4: _____ by plane

Figure 1: An overview of Cross-View Training. The model is trained with standard supervised learning on labeled examples. On unlabeled examples, auxiliary prediction modules with different views of the input are trained to agree with the primary prediction module. This particular example shows CVT applied to named entity recognition. From the labeled example, the model can learn that “Washington” usually refers to a location. Then, on unlabeled data, auxiliary prediction modules are trained to reach the same prediction without seeing some of the input. In doing so, they improve the contextual representations produced by the model, for example, learning that “traveled to” is usually followed by a location.

duced by the model with parameters θ on input x_i . During CVT, the model alternates learning on a minibatch of labeled examples and learning on a minibatch of unlabeled examples. For labeled examples, CVT uses standard cross-entropy loss:

$$\mathcal{L}_{\text{sup}}(\theta) = \frac{1}{|\mathcal{D}_l|} \sum_{x_i, y_i \in \mathcal{D}_l} CE(y_i, p_\theta(y|x_i))$$

CVT adds k auxiliary prediction modules to the model, which are used when learning on unlabeled examples. A prediction module is usually a small neural network (e.g., a hidden layer followed by a softmax layer). Each one takes as input an intermediate representation $h^j(x_i)$ produced by the model (e.g., the outputs of one of the LSTMs in a Bi-LSTM model). It outputs a distribution over labels $p_\theta^j(y|x_i)$. Each h^j is chosen such that it only uses a part of the input x_i ; the particular choice

can depend on the task and model architecture. We propose variants for several tasks in Section 3. The auxiliary prediction modules are only used during training; the test-time prediction come from the primary prediction module that produces p_θ .

On an unlabeled example, the model first produces soft targets $p_\theta(y|x_i)$ by performing inference. CVT trains the auxiliary prediction modules to match the primary prediction module on the unlabeled data by minimizing

$$\mathcal{L}_{\text{CVT}}(\theta) = \frac{1}{|\mathcal{D}_{ul}|} \sum_{x_i \in \mathcal{D}_{ul}} \sum_{j=1}^k D(p_\theta(y|x_i), p_\theta^j(y|x_i))$$

where D is a distance function between probability distributions (we use KL divergence). We hold the primary module’s prediction $p_\theta(y|x_i)$ fixed during training (i.e., we do not back-propagate through it) so the auxiliary modules learn to imitate the primary one, but not vice versa. CVT works by enhancing the model’s representation learning. As the auxiliary modules train, the representations they take as input improve so they are useful for making predictions even when some of the model’s inputs are not available. This in turn improves the primary prediction module, which is built on top of the same shared representations.

We combine the supervised and CVT losses into the total loss, $\mathcal{L} = \mathcal{L}_{\text{sup}} + \mathcal{L}_{\text{CVT}}$, and minimize it with stochastic gradient descent. In particular, we alternate minimizing \mathcal{L}_{sup} over a minibatch of labeled examples and minimizing \mathcal{L}_{CVT} over a minibatch of unlabeled examples.

For most neural networks, adding a few additional prediction modules is computationally cheap compared to the portion of the model building up representations (such as an RNN or CNN). Therefore our method contributes little overhead to training time over other self-training approaches for most tasks. CVT does not change inference time or the number of parameters in the fully-trained model because the auxiliary prediction modules are only used during training.

2.2 Combining CVT with Multi-Task Learning

CVT can easily be combined with multi-task learning by adding additional prediction modules for the other tasks on top of the shared Bi-LSTM encoder. During supervised learning, we randomly select a task and then update \mathcal{L}_{sup} using a minibatch of labeled data for that task. When learning on the unlabeled data, we optimize \mathcal{L}_{CVT}

jointly across all tasks at once, first running inference with all the primary prediction modules and then learning from the predictions with all the auxiliary prediction modules. As before, the model alternates training on minibatches of labeled and unlabeled examples.

Examples labeled across many tasks are useful for multi-task systems to learn from, but most datasets are only labeled with one task. A benefit of multi-task CVT is that the model creates (artificial) all-tasks-labeled examples from unlabeled data. This significantly improves the model’s data efficiency and training time. Since running prediction modules is computationally cheap, computing \mathcal{L}_{CVT} is not much slower for many tasks than it is for a single one. However, we find the all-tasks-labeled examples substantially speed up model convergence. For example, our model trained on six tasks takes about three times as long to converge as the average model trained on one task, a 2x decrease in total training time.

3 Cross-View Training Models

CVT relies on auxiliary prediction modules that have restricted views of the input. In this section, we describe specific constructions of the auxiliary prediction modules that are effective for sequence tagging, dependency parsing, and sequence-to-sequence learning.

3.1 Bi-LSTM Sentence Encoder

All of our models use a two-layer CNN-BiLSTM (Chiu and Nichols, 2016; Ma and Hovy, 2016) sentence encoder. It takes as input a sequence of words $x_i = [x_i^1, x_i^2, \dots, x_i^T]$. First, each word is represented as the sum of an embedding vector and the output of a character-level Convolutional Neural Network, resulting in a sequence of vectors $v = [v^1, v^2, \dots, v^T]$. The encoder applies a two-layer bidirectional LSTM (Graves and Schmidhuber, 2005) to these representations. The first layer runs a Long Short-Term Memory unit (Hochreiter and Schmidhuber, 1997) in the forward direction (taking v^t as input at each step t) and the backward direction (taking v^{T-t+1} at each step) to produce vector sequences $[\vec{h}_1^1, \vec{h}_1^2, \dots, \vec{h}_1^T]$ and $[\overleftarrow{h}_1^1, \overleftarrow{h}_1^2, \dots, \overleftarrow{h}_1^T]$. The output of the Bi-LSTM is the concatenation of these vectors: $h_1 = [\vec{h}_1^1 \oplus \overleftarrow{h}_1^1, \dots, \vec{h}_1^T \oplus \overleftarrow{h}_1^T]$. The second Bi-LSTM layer works the same, producing outputs h_2 , except it takes h_1 as input instead of v .

3.2 CVT for Sequence Tagging

In sequence tagging, each token x_i^t has a corresponding label y_i^t . The primary prediction module for sequence tagging produces a probability distribution over classes for the t^{th} label using a one-hidden-layer neural network applied to the corresponding encoder outputs:

$$\begin{aligned} p(y^t|x_i) &= \text{NN}(h_1^t \oplus h_2^t) \\ &= \text{softmax}(U \cdot \text{ReLU}(W(h_1^t \oplus h_2^t)) + b) \end{aligned}$$

The auxiliary prediction modules take $\vec{h}_1(x_i)$ and $\overleftarrow{h}_1(x_i)$, the outputs of the forward and backward LSTMs in the first² Bi-LSTM layer, as inputs. We add the following four auxiliary prediction modules to the model (see Figure 2):

$$\begin{aligned} p_\theta^{\text{fwd}}(y^t|x_i) &= \text{NN}^{\text{fwd}}(\vec{h}_1^t(x_i)) \\ p_\theta^{\text{bwd}}(y^t|x_i) &= \text{NN}^{\text{bwd}}(\overleftarrow{h}_1^t(x_i)) \\ p_\theta^{\text{future}}(y^t|x_i) &= \text{NN}^{\text{future}}(\vec{h}_1^{t-1}(x_i)) \\ p_\theta^{\text{past}}(y^t|x_i) &= \text{NN}^{\text{past}}(\overleftarrow{h}_1^{t+1}(x_i)) \end{aligned}$$

The ‘‘forward’’ module makes each prediction without seeing the right context of the current token. The ‘‘future’’ module makes each prediction without the right context or the current token itself. Therefore it works like a neural language model that, instead of predicting which token comes next in the sequence, predicts which class of token comes next. The ‘‘backward’’ and ‘‘past’’ modules are analogous.

3.3 CVT for Dependency Parsing

In a dependency parse, words in a sentence are treated as nodes in a graph. Typed directed edges connect the words, forming a tree structure describing the syntactic structure of the sentence. In particular, each word x_i^t in a sentence $x_i = x_i^1, \dots, x_i^T$ receives exactly one in-going edge (u, t, r) going from word x_i^u (called the ‘‘head’’) to it (the ‘‘dependent’’) of type r (the ‘‘relation’’). We use a graph-based dependency parser similar to the one from [Dozat and Manning \(2017\)](#). This treats dependency parsing as a classification task where the goal is to predict which in-going edge $y_i^t = (u, t, r)$ connects to each word x_i^t .

First, the representations produced by the encoder for the candidate head and dependent are

²Modules taking inputs from the second Bi-LSTM layer would not have restricted views because information about the whole sentence gets propagated through the first layer.




Figure 2: Auxiliary prediction modules for sequence tagging models. Each one sees a restricted view of the input. For example, the ‘‘forward’’ prediction module does not see any context to the right of the current token when predicting that token’s label. For simplicity, we only show a one layer Bi-LSTM encoder and only show the model’s predictions for a single time step.

passed through separate hidden layers. A bilinear classifier applied to these representations produces a score for each candidate edge. Lastly, these scores are passed through a softmax layer to produce probabilities. Mathematically, the probability of an edge is given as:

$$p_\theta((u, t, r)|x_i) \propto e^{s(h_1^u(x_i) \oplus h_2^u(x_i), h_1^t(x_i) \oplus h_2^t(x_i), r)}$$

where s is the scoring function:

$$\begin{aligned} s(z_1, z_2, r) &= \text{ReLU}(W_{\text{head}}z_1 + b_{\text{head}})(W_r + W) \\ &\quad \text{ReLU}(W_{\text{dep}}z_2 + b_{\text{dep}}) \end{aligned}$$

The bilinear classifier uses a weight matrix W_r specific to the candidate relation as well as a weight matrix W shared across all relations. Note that unlike in most prior work, our dependency parser only takes words as inputs, not words and part-of-speech tags.

We add four auxiliary prediction modules to our model for cross-view training:

$$\begin{aligned} p_\theta^{\text{fwd-fwd}}((u, t, r)|x_i) &\propto e^{s^{\text{fwd-fwd}}(\vec{h}_1^u(x_i), \vec{h}_1^t(x_i), r)} \\ p_\theta^{\text{fwd-bwd}}((u, t, r)|x_i) &\propto e^{s^{\text{fwd-bwd}}(\vec{h}_1^u(x_i), \overleftarrow{h}_1^t(x_i), r)} \\ p_\theta^{\text{bwd-fwd}}((u, t, r)|x_i) &\propto e^{s^{\text{bwd-fwd}}(\overleftarrow{h}_1^u(x_i), \vec{h}_1^t(x_i), r)} \\ p_\theta^{\text{bwd-bwd}}((u, t, r)|x_i) &\propto e^{s^{\text{bwd-bwd}}(\overleftarrow{h}_1^u(x_i), \overleftarrow{h}_1^t(x_i), r)} \end{aligned}$$

Each one has some missing context (not seeing either the preceding or following words) for the candidate head and candidate dependent.

3.4 CVT for Sequence-to-Sequence Learning

We use an encoder-decoder sequence-to-sequence model with attention (Sutskever et al., 2014; Bahdanau et al., 2015). Each example consists of an input (source) sequence $x_i = x_i^1, \dots, x_i^T$ and output (target) sequence $y_i = y_i^1, \dots, y_i^K$. The encoder’s representations are passed into an LSTM decoder using a bilinear attention mechanism (Luong et al., 2015). In particular, at each time step t the decoder computes an attention distribution over source sequence hidden states as $\alpha_j \propto e^{h^j W_a \bar{h}^t}$ where \bar{h}^t is the decoder’s current hidden state. The source hidden states weighted by the attention distribution form a context vector: $c_t = \sum_j \alpha_j h^j$. Next, the context vector and current hidden state are combined into an attention vector $a_t = \tanh(W_a[c_t, h_t])$. Lastly, a softmax layer predicts the next token in the output sequence: $p(y_i^t | y_i^{<t}, x_i) = \text{softmax}(W_s a_t)$.

We add two auxiliary decoders when applying CVT. The auxiliary decoders share embedding and LSTM parameters with the primary decoder, but have different parameters for the attention mechanisms and softmax layers. For the first one, we restrict its view of the input by applying attention dropout, randomly zeroing out a fraction of its attention weights. The second one is trained to predict the next word in the target sequence rather than the current one: $p_\theta^{\text{future}}(y_i^t | y_i^{<t}, x_i) = \text{softmax}(W_s^{\text{future}} a_{t-1}^{\text{future}})$. Since there is no target sequence for unlabeled examples, we cannot apply teacher forcing to get an output distribution over the vocabulary from the primary decoder at each time step. Instead, we produce hard targets for the auxiliary modules by running the primary decoder with beam search on the input sequence. This idea has previously been applied to sequence-level knowledge distillation by Kim and Rush (2016).

4 Experiments

We compare Cross-View Training against several strong baselines on seven tasks:

Combinatory Categorical Grammar (CCG) Supertagging: We use data from CCGBank (Hockenmaier and Steedman, 2007).

Text Chunking: We use the CoNLL-2000 data (Tjong Kim Sang and Buchholz, 2000).

Named Entity Recognition (NER): We use the CoNLL-2003 data (Tjong Kim Sang and De Meulder, 2003).

Fine-Grained NER (FGN): We use the OntoNotes (Hovy et al., 2006) dataset.

Part-of-Speech (POS) Tagging: We use the Wall Street Journal portion of the Penn Treebank (Marcus et al., 1993).

Dependency Parsing: We use the Penn Treebank converted to Stanford Dependencies version 3.3.0.

Machine Translation: We use the English-Vietnamese translation dataset from IWSLT 2015 (Cettolo et al., 2015). We report (tokenized) BLEU scores on the tst2013 test set.

We use the 1 Billion Word Language Model Benchmark (Chelba et al., 2014) as a pool of unlabeled sentences for semi-supervised learning.

4.1 Model Details and Baselines

We apply dropout during training, but not when running the primary prediction module to produce soft targets on unlabeled examples. In addition to the auxiliary prediction modules listed in Section 3, we find it slightly improves results to add another one that sees the whole input rather than a subset (but unlike the primary prediction module, does have dropout applied to its representations). Unless indicated otherwise, our models have LSTMs with 1024-sized hidden states and 512-sized projection layers. See the supplementary material for full training details and hyperparameters. We compare CVT with the following other semi-supervised learning algorithms:

Word Dropout. In this method, we only train the primary prediction module. When acting as a teacher it is run as normal, but when acting as a student, we randomly replace some of the input words with a REMOVED token. This is similar to CVT in that it exposes the model to a restricted view of the input. However, it is less data efficient. By carefully designing the auxiliary prediction modules, it is possible to train the auxiliary prediction modules to match the primary one across many different views of the input a once, rather than just one view at a time.

Virtual Adversarial Training (VAT). VAT (Miyato et al., 2016) works like word dropout, but adds noise to the word embeddings of the student instead of dropping out words. Notably, the noise is chosen adversarially so it most changes the model’s prediction. This method was applied successfully to semi-supervised text classification

Method	CCG Acc.	Chunk F1	NER F1	FGN F1	POS Acc.	Dep. UAS	Parse LAS	Translate BLEU
Shortcut LSTM (Wu et al., 2017)	95.1				97.53			
ID-CNN-CRF (Strubell et al., 2017)			90.7	86.8				
JMT [†] (Hashimoto et al., 2017)		95.8			97.55	94.7	92.9	
TagLM* (Peters et al., 2017)		96.4	91.9					
ELMo* (Peters et al., 2018)			92.2					
Biaffine (Dozat and Manning, 2017)						95.7	94.1	
Stack Pointer (Ma et al., 2018)						95.9	94.2	
Stanford (Luong and Manning, 2015)								23.3
Google (Luong et al., 2017)								26.1
Supervised	94.9	95.1	91.2	87.5	97.60	95.1	93.3	28.9
Virtual Adversarial Training*	95.1	95.1	91.8	87.9	97.64	95.4	93.7	–
Word Dropout*	95.2	95.8	92.1	88.1	97.66	95.6	93.8	29.3
ELMo (our implementation)*	95.8	96.5	92.2	88.5	97.72	96.2	94.4	29.3
ELMo + Multi-task* [†]	95.9	96.8	92.3	88.4	97.79	96.4	94.8	–
CVT*	95.7	96.6	92.3	88.7	97.70	95.9	94.1	29.6
CVT + Multi-task* [†]	96.0	96.9	92.4	88.4	97.76	96.4	94.8	–
CVT + Multi-task + Large* [†]	96.1	97.0	92.6	88.8	97.74	96.6	95.0	–

Table 1: Results on the test sets. We report the mean score over 5 runs. Standard deviations in score are around 0.1 for NER, FGN, and translation, 0.02 for POS, and 0.05 for the other tasks. See the supplementary materials for results with them included. The +Large model has four times as many hidden units as the others, making it similar in size to the models when ELMo is included. * denotes semi-supervised and [†] denotes multi-task.

by Miyato et al. (2017).

ELMo. ELMo incorporates the representations from a large separately-trained language model into a task-specific model. Our implementation follows Peters et al. (2018). When combining ELMo with multi-task learning, we allow each task to learn its own weights for the ELMo embeddings going into each prediction module. We found applying dropout to the ELMo embeddings was crucial for achieving good performance.

4.2 Results

Results are shown in Table 1. CVT on its own outperforms or is comparable to the best previously published results on all tasks. Figure 3 shows an example win for CVT over supervised learning.

Of the prior results listed in Table 1, only TagLM and ELMo are semi-supervised. These methods first train an enormous language model on unlabeled data and incorporate the representations produced by the language model into a supervised classifier. Our base models use 1024 hidden units in their LSTMs (compared to 4096 in ELMo), require fewer training steps (around one pass over the billion-word benchmark rather than

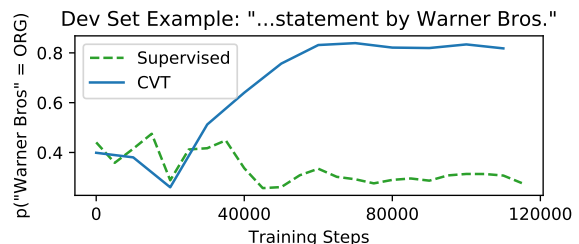


Figure 3: An NER example that CVT classifies correctly but supervised learning does not. “Warner” only occurs as a last name in the train set, so the supervised model classifies “Warner Bros” as a person. The CVT model also mistakenly classifies “Warner Bros” as a person to start with, but as it sees more of the unlabeled data (in which “Warner” occurs thousands of times) it learns that “Warner Bros” is an organization.

many passes), and do not require a pipelined training procedure. Therefore, although they perform on par with ELMo, they are faster and simpler to train. Increasing the size of our CVT+Multi-task model so it has 4096 units in its LSTMs like ELMo improves results further so they are significantly better than the ELMo+Multi-task ones. We suspect there could be further gains from combining our method with language model pre-training, which we leave for future work.

CVT + Multi-Task. We train a single shared-encoder CVT model to perform all of the tasks except machine translation (as it is quite different and requires more training time than the other ones). Multi-task learning improves results on all of the tasks except fine-grained NER, sometimes by large margins. Prior work on many-task NLP such as Hashimoto et al. (2017) uses complicated architectures and training algorithms. Our result shows that simple parameter sharing can be enough for effective many-task learning when the model is big and trained on a large amount of data.

Interestingly, multi-task learning works better in conjunction with CVT than with ELMo. We hypothesize that the ELMo models quickly fit to the data primarily using the ELMo vectors, which perhaps hinders the model from learning effective representations that transfer across tasks. We also believe CVT alleviates the danger of the model “forgetting” one task while training on the other ones, a well-known problem in many-task learning (Kirkpatrick et al., 2017). During multi-task CVT, the model makes predictions about unlabeled examples across all tasks, creating (artificial) all-tasks-labeled examples, so the model does not only see one task at a time. In fact, multi-task learning plus self training is similar to the Learning without Forgetting algorithm (Li and Hoiem, 2016), which trains the model to keep its predictions on an old task unchanged when learning a new task. To test the value of all-tasks-labeled examples, we trained a multi-task CVT model that only computes \mathcal{L}_{CVT} on one task at a time (chosen randomly for each unlabeled minibatch) instead of for all tasks in parallel. The one-at-a-time model performs substantially worse (see Table 2).

Model	CCG	Chnk	NER	FGN	POS	Dep.
CVT-MT	95.7	97.4	96.0	86.7	97.74	94.4
w/out parallel	95.4	97.1	95.6	86.3	97.71	94.1

Table 2: Dev set performance of multi-task CVT with and without producing all-tasks-labeled examples.

Model Generalization. In order to evaluate how our models generalize to the dev set from the train set, we plot the dev vs. train accuracy for our different methods as they learn (see Figure 4). Both CVT and multi-task learning improve model generalization: for the same train accuracy, the models get better dev accuracy than purely supervised learning. Interestingly, CVT continues to improve

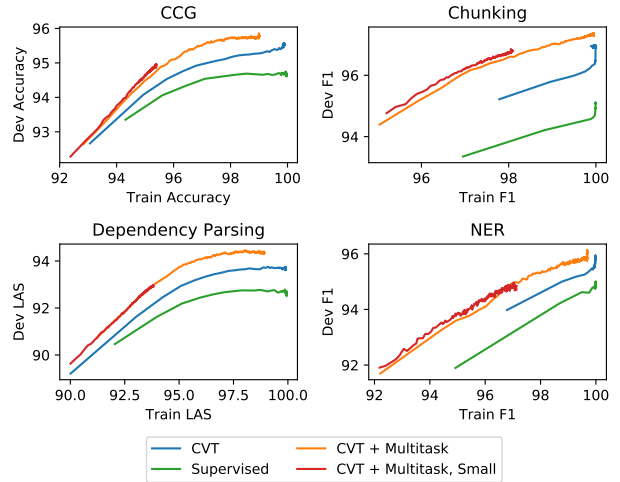


Figure 4: Dev set vs. Train set accuracy for various methods. The “small” model has 1/4 the LSTM hidden state size of the other ones (256 instead of 1024).

in dev set accuracy while close to 100% train accuracy for CCG, Chunking, and NER, perhaps because the model is still learning from unlabeled data even when it has completely fit to the train set. We also show results for a smaller multi-task + CVT model. Although it generalizes at least as well as the larger one, it halts making progress on the train set earlier. This suggests it is important to use sufficiently large neural networks for multi-task learning: otherwise the model does not have the capacity to fit to all the training data.

Auxiliary Prediction Module Ablation. We briefly explore which auxiliary prediction modules are more important for the sequence tagging tasks in Table 3. We find that both kinds of auxiliary prediction modules improve performance, but that the future and past modules improve results more than the forward and backward ones, perhaps because they see a more restricted and challenging view of the input.

Model	CCG	Chnk	NER	FGN	POS
Supervised	94.8	95.5	95.0	86.0	97.59
CVT	95.6	97.0	95.9	87.3	97.66
no fwd/bwd	-0.1	-0.2	-0.2	-0.1	-0.01
no future/past	-0.3	-0.4	-0.4	-0.3	-0.04

Table 3: Ablation study on auxiliary prediction modules for sequence tagging.

Training Models on Small Datasets. We explore how CVT scales with dataset size by varying the amount of training data the model has ac-

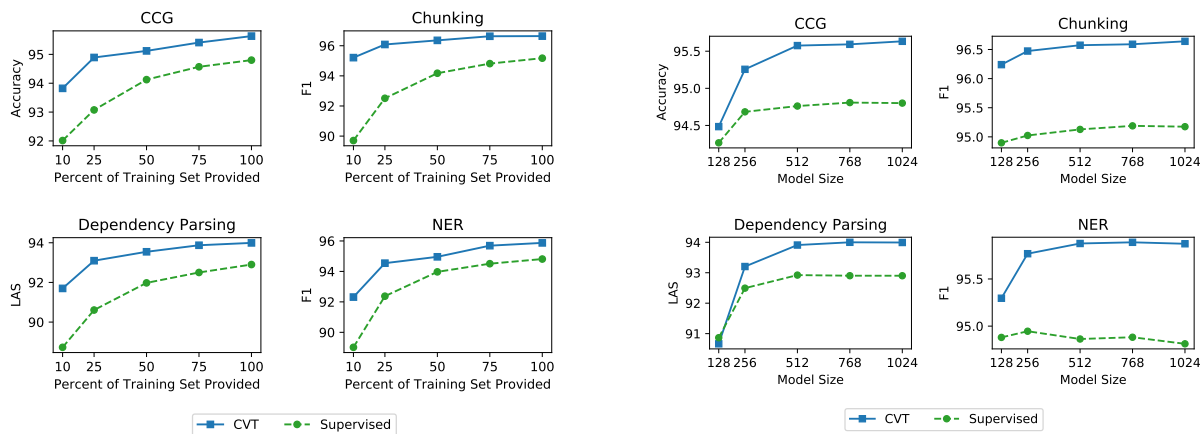


Figure 5: Left: Dev set performance vs. percent of the training set provided to the model. Right: Dev set performance vs. model size. The x axis shows the number of hidden units in the LSTM layers; the projection layers and other hidden layers in the network are half that size. Points correspond to the mean of three runs.

cess to. Unsurprisingly, the improvement of CVT over purely supervised learning grows larger as the amount of labeled data decreases (see Figure 5, left). Using only 25% of the labeled data, our approach already performs as well or better than a fully supervised model using 100% of the training data, demonstrating that CVT is particularly useful on small datasets.

Training Larger Models. Most sequence taggers and dependency parsers in prior work use small LSTMs (hidden state sizes of around 300) because larger models yield little to no gains in performance (Reimers and Gurevych, 2017). We found our own supervised approaches also do not benefit greatly from increasing the model size. In contrast, when using CVT accuracy scales better with model size (see Figure 5, right). This finding suggests the appropriate semi-supervised learning methods may enable the development of larger, more sophisticated models for NLP tasks with limited amounts of labeled data.

Generalizable Representations. Lastly, we explore training the CVT+multi-task model on five tasks, freezing the encoder, and then only training a prediction module on the sixth task. This tests whether the encoder’s representations generalize to a new task not seen during its training. Only training the prediction module is very fast because (1) the encoder (which is by far the slowest part of the model) has to be run over each example only once and (2) we do not back-propagate into the encoder. Results are shown in Table 4.

Training only a prediction module on top of multi-task representations works remarkably well,

Model	CCG	Chnk	NER	FGN	POS	Dep.
Supervised	94.8	95.6	95.0	86.0	97.59	92.9
CVT-MT frozen	95.1	96.6	94.6	83.2	97.66	92.5
ELMo frozen	94.3	92.2	91.3	80.6	97.50	89.4

Table 4: Comparison of single-task models on the dev sets. “CVT-MT frozen” means we pretrain a CVT + multi-task model on five tasks, and then train only the prediction module for the sixth. “ELMo frozen” means we train prediction modules (but no LSTMs) on top of ELMo embeddings.

outperforming ELMo embeddings and sometimes even a vanilla supervised model, showing the multi-task model is building up effective representations for language. In particular, the representations could be used like skip-thought vectors (Kiros et al., 2015) to quickly train models on new tasks without slow representation learning.

5 Related Work

Unsupervised Representation Learning. Early approaches to deep semi-supervised learning pre-train neural models on unlabeled data, which has been successful for applications in computer vision (Jarrett et al., 2009; LeCun et al., 2010) and NLP. Particularly noteworthy for NLP are algorithms for learning effective word embeddings (Collobert et al., 2011; Mikolov et al., 2013; Pennington et al., 2014) and language model pre-training (Dai and Le, 2015; Ramachandran et al., 2017; Peters et al., 2018; Howard and Ruder, 2018; Radford et al., 2018). Pre-training on other tasks such as machine translation has also been studied (McCann et al., 2017). Other approaches train

“thought vectors” representing sentences through unsupervised (Kiros et al., 2015; Hill et al., 2016) or supervised (Conneau et al., 2017) learning.

Self-Training. One of the earliest approaches to semi-supervised learning is self-training (Scudder, 1965), which has been successfully applied to NLP tasks such as word-sense disambiguation (Yarowsky, 1995) and parsing (McClosky et al., 2006). In each round of training, the classifier, acting as a “teacher,” labels some of the unlabeled data and adds it to the training set. Then, acting as a “student,” it is retrained on the new training set. Many recent approaches (including the consistency regularization methods discussed below and our own method) train the student with soft targets from the teacher’s output distribution rather than a hard label, making the procedure more akin to knowledge distillation (Hinton et al., 2015). It is also possible to use multiple models or prediction modules for the teacher, such as in tri-training (Zhou and Li, 2005; Ruder and Plank, 2018).

Consistency Regularization. Recent works add noise (e.g., drawn from a Gaussian distribution) or apply stochastic transformations (e.g., horizontally flipping an image) to the student’s inputs. This trains the model to give consistent predictions to nearby data points, encouraging distributional smoothness in the model. Consistency regularization has been very successful for computer vision applications (Bachman et al., 2014; Laine and Aila, 2017; Tarvainen and Valpola, 2017). However, stochastic input alterations are more difficult to apply to discrete data like text, making consistency regularization less used for natural language processing. One solution is to add noise to the model’s word embeddings (Miyato et al., 2017); we compare against this approach in our experiments. CVT is easily applicable to text because it does not require changing the student’s inputs.

Multi-View Learning. Multi-view learning on data where features can be separated into distinct subsets has been well studied (Xu et al., 2013). Particularly relevant are co-training (Blum and Mitchell, 1998) and co-regularization (Sindhwani and Belkin, 2005), which trains two models with disjoint views of the input. On unlabeled data, each one acts as a “teacher” for the other model. In contrast to these methods, our approach trains a single unified model where auxiliary prediction modules see different, but not necessarily indepen-

dent views of the input.

Self Supervision. Self-supervised learning methods train auxiliary prediction modules on tasks where performance can be measured without human-provided labels. Recent work has jointly trained image classifiers with tasks like relative position and colorization (Doersch and Zisserman, 2017), sequence taggers with language modeling (Rei, 2017), and reinforcement learning agents with predicting changes in the environment (Jaderberg et al., 2017). Unlike these approaches, our auxiliary losses are based on self-labeling, not labels deterministically constructed from the input.

Multi-Task Learning. There has been extensive prior work on multi-task learning (Caruana, 1997; Ruder, 2017). For NLP, most work has focused on a small number of closely related tasks (Luong et al., 2016; Zhang and Weiss, 2016; Søgaard and Goldberg, 2016; Peng et al., 2017). Many-task systems are less commonly developed. Collobert and Weston (2008) propose a many-task system sharing word embeddings between the tasks, Hashimoto et al. (2017) train a many-task model where the tasks are arranged hierarchically according to their linguistic level, and Subramanian et al. (2018) train a shared-encoder many-task model for the purpose of learning better sentence representations for use in downstream tasks, not for improving results on the original tasks.

6 Conclusion

We propose Cross-View Training, a new method for semi-supervised learning. Our approach allows models to effectively leverage their own predictions on unlabeled data, training them to produce effective representations that yield accurate predictions even when some of the input is not available. We achieve excellent results across seven NLP tasks, especially when CVT is combined with multi-task learning.

Acknowledgements

We thank Abi See, Christopher Clark, He He, Peng Qi, Reid Pryzant, Yuaho Zhang, and the anonymous reviewers for their thoughtful comments and suggestions. We thank Takeru Miyato for help with his virtual adversarial training code and Emma Strubell for answering our questions about OntoNotes NER. Kevin is supported by a Google PhD Fellowship.

References

- Philip Bachman, Ouais Alsharif, and Doina Precup. 2014. Learning with pseudo-ensembles. In *NIPS*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Avrim Blum and Tom Mitchell. 1998. Combining labeled and unlabeled data with co-training. In *COLT*. ACM.
- Rich Caruana. 1997. Multitask learning. *Machine Learning*, 28:41–75.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2015. The IWSLT 2015 evaluation campaign. In *International Workshop on Spoken Language Translation*.
- Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn, and Tony Robinson. 2014. One billion word benchmark for measuring progress in statistical language modeling. In *INTERSPEECH*.
- Jason PC Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional LSTM-CNNs. *Transactions of the Association for Computational Linguistics*.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: deep neural networks with multitask learning. In *ICML*.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *EMNLP*.
- Andrew M Dai and Quoc V Le. 2015. Semi-supervised sequence learning. In *NIPS*.
- Carl Doersch and Andrew Zisserman. 2017. Multi-task self-supervised visual learning. *arXiv preprint arXiv:1708.07860*.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *ICLR*.
- Alex Graves and Jürgen Schmidhuber. 2005. Frame-wise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks*, 18(5):602–610.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2017. A joint many-task model: Growing a neural network for multiple nlp tasks. In *EMNLP*.
- Felix Hill, Kyunghyun Cho, and Anna Korhonen. 2016. Learning distributed representations of sentences from unlabelled data. In *HLT-NAACL*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Julia Hockenmaier and Mark Steedman. 2007. CCG-bank: a corpus of CCG derivations and dependency structures extracted from the Penn treebank. *Computational Linguistics*, 33(3):355–396.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. Ontonotes: the 90% solution. In *HLT-NAACL*.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *ACL*.
- Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. 2017. Reinforcement learning with unsupervised auxiliary tasks. In *ICLR*.
- Kevin Jarrett, Koray Kavukcuoglu, Yann LeCun, et al. 2009. What is the best multi-stage architecture for object recognition? In *IEEE Conference on Computer Vision*.
- Yoon Kim and Alexander M. Rush. 2016. Sequence-level knowledge distillation. In *EMNLP*.
- James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114 13:3521–3526.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- Samuli Laine and Timo Aila. 2017. Temporal ensembling for semi-supervised learning. In *ICLR*.
- Yann LeCun, Koray Kavukcuoglu, and Clément Faret. 2010. Convolutional networks and applications in vision. In *ISCAS*. IEEE.
- Zhizhong Li and Derek Hoiem. 2016. Learning without forgetting. In *ECCV*.
- Minh-Thang Luong, Eugene Brevdo, and Rui Zhao. 2017. Neural machine translation (seq2seq) tutorial. <https://github.com/tensorflow/nmt>.

- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-task sequence to sequence learning. In *ICLR*.
- Minh-Thang Luong and Christopher D. Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *IWSLT*.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional LSTM-CNN-CRF. In *ACL*.
- Xuezhe Ma, Zecong Hu, Jingzhou Liu, Nanyun Peng, Graham Neubig, and Eduard Hovy. 2018. Stack-pointer networks for dependency parsing. In *ACL*.
- Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The Penn treebank. *Computational linguistics*, 19(2):313–330.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. Learned in translation: Contextualized word vectors. In *NIPS*.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. Effective self-training for parsing. In *ACL*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- Takeru Miyato, Andrew M Dai, and Ian Goodfellow. 2017. Adversarial training methods for semi-supervised text classification. In *ICLR*.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. 2016. Distributional smoothing with virtual adversarial training. In *ICLR*.
- Hao Peng, Sam Thomson, and Noah A. Smith. 2017. Deep multitask learning for semantic dependency parsing. In *ACL*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Matthew E Peters, Waleed Ammar, Chandra Bhagavatula, and Russell Power. 2017. Semi-supervised sequence tagging with bidirectional language models. In *ACL*.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. <https://blog.openai.com/language-unsupervised>.
- Prajit Ramachandran, Peter J Liu, and Quoc V Le. 2017. Unsupervised pretraining for sequence to sequence learning. In *EMNLP*.
- Marek Rei. 2017. Semi-supervised multitask learning for sequence labeling. In *ACL*.
- Nils Reimers and Iryna Gurevych. 2017. Reporting score distributions makes a difference: Performance study of LSTM-networks for sequence tagging. In *EMNLP*.
- Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Sebastian Ruder and Barbara Plank. 2018. Strong baselines for neural semi-supervised learning under domain shift. In *ACL*.
- Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. 2016. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NIPS*.
- H Scudder. 1965. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371.
- Vikas Sindhwani and Mikhail Belkin. 2005. A co-regularization approach to semi-supervised learning with multiple views. In *ICML Workshop on Learning with Multiple Views*.
- Anders Søgaard and Yoav Goldberg. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *ACL*.
- Emma Strubell, Patrick Verga, David Belanger, and Andrew McCallum. 2017. Fast and accurate sequence labeling with iterated dilated convolutions. In *EMNLP*.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. Learning general purpose distributed sentence representations via large scale multi-task learning. In *ICLR*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- Antti Tarvainen and Harri Valpola. 2017. Weight-averaged consistency targets improve semi-supervised deep learning results. In *Workshop on Learning with Limited Labeled Data, NIPS*.
- Erik F Tjong Kim Sang and Sabine Buchholz. 2000. Introduction to the CoNLL-2000 shared task: Chunking. In *CoNLL*.

- Erik F Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *HLT-NAACL*.
- Xiang Wei, Zixia Liu, Liqiang Wang, and Boqing Gong. 2018. Improving the improved training of Wasserstein GANs. In *ICLR*.
- Huijia Wu, Jiajun Zhang, and Chengqing Zong. 2017. Shortcut sequence tagging. *arXiv preprint arXiv:1701.00576*.
- Chang Xu, Dacheng Tao, and Chao Xu. 2013. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*.
- David Yarowsky. 1995. Unsupervised word sense disambiguation rivaling supervised methods. In *ACL*.
- Yuan Zhang and David Weiss. 2016. Stack-propagation: Improved representation learning for syntax. In *ACL*.
- Zhi-Hua Zhou and Ming Li. 2005. Tri-training: Exploiting unlabeled data using three classifiers. *IEEE Transactions on knowledge and Data Engineering*.