

Self-Supervised Self-Supervision by Combining Deep Learning and Probabilistic Logic

Hunter Lang,^{1*} Hoifung Poon²

MIT¹

Microsoft Research²

hjl@mit.edu, hoifung@microsoft.com

Abstract

Labeling training examples at scale is a perennial challenge in machine learning. Self-supervision methods compensate for the lack of direct supervision by leveraging prior knowledge to automatically generate noisy labeled examples. Deep probabilistic logic (DPL) is a unifying framework for self-supervised learning that represents unknown labels as latent variables and incorporates diverse self-supervision using probabilistic logic to train a deep neural network end-to-end using variational EM. While DPL is successful at combining pre-specified self-supervision, manually crafting self-supervision to attain high accuracy may still be tedious and challenging. In this paper, we propose Self-Supervised Self-Supervision (S4), which adds to DPL the capability to learn new self-supervision automatically. Starting from an initial “seed,” S4 iteratively uses the deep neural network to propose new self supervision. These are either added directly (a form of *structured* self-training) or verified by a human expert (as in feature-based active learning). Experiments show that S4 is able to automatically propose accurate self-supervision and can often nearly match the accuracy of supervised methods with a tiny fraction of the human effort.

Introduction

Machine learning has made great strides in enhancing model sophistication and learning efficacy, as exemplified by recent advances in deep learning (LeCun, Bengio, and Hinton 2015). However, contemporary supervised learning techniques require a large amount of labeled data, which is expensive and time-consuming to produce. This problem is particularly acute in specialized domains like biomedicine, where crowdsourcing is difficult to apply. Self-supervision has emerged as a promising paradigm to overcome the annotation bottleneck by automatically generating noisy training examples from unlabeled data. In particular, *task-specific self-supervision* converts prior knowledge into self-supervision templates for label generation, as in distant supervision (Mintz et al. 2009), data programming (Ratner et al. 2016), and joint inference (Poon and Domingos 2008).

Deep probabilistic logic (DPL) is a unifying framework for self-supervision that combines deep learning with probabilistic logic (Wang and Poon 2018). It represents unknown

labels as latent variables, and incorporates prior beliefs over labels and their dependencies as virtual evidence in a graphical model. The marginal beliefs over the latent variables are used as probabilistic labels to train a deep neural network for the end prediction task. The trained neural network in turn provides belief updates to refine the graphical model parameters, and the process continues, using variational EM.

Wang and Poon (2018) show that DPL can effectively combine diverse sources of self-supervision in a coherent probabilistic framework and subsume supervised and semi-supervised learning as special cases. While promising, DPL and related approaches still require human experts to manually specify self-supervision. This is particularly challenging for self-supervision such as data programming and joint inference, which require domain expertise and extensive effort to identify the many relevant virtual evidences for attaining high accuracy in the end task.

In this paper, we propose *Self-Supervised Self-Supervision (S4)* as a general framework for learning to add new self-supervision. In particular, we extend deep probabilistic logic (DPL) with structure learning and active learning capabilities (see Figure 1). After running DPL using the pre-specified seed self-supervision, S4 iteratively proposes new virtual evidence using the trained deep neural network and graphical model, and determines whether to add this evidence directly to the graphical model or ask a human expert to vet it. The former can be viewed as *structured self-training*, which generalizes self-training (e.g., McClosky, Charniak, and Johnson 2006) by adding not only individual labels but also arbitrary probabilistic factors over them. The latter subsumes *feature-based active learning* (Druck, Settles, and McCallum 2009) with arbitrary features expressible using probabilistic logic. By combining the two in a unified framework, S4 can leverage both paradigms for generating new self-supervision and subsume many related approaches as special cases.

We use transformer-based models for the deep neural network in DPL and explore various self-supervision proposal mechanisms based on neural attention and label entropy. Our method can learn to propose both unary potential factors over individual labels and joint-inference factors over multiple labels. We conducted experiments on various natural language processing (NLP) tasks to explore the potential of our method. We held out gold labels for evaluation

*Work done while at Microsoft Research.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

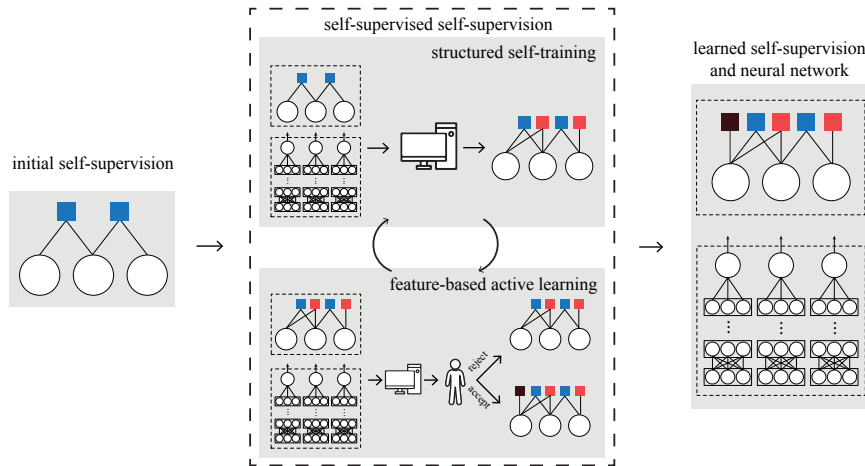


Figure 1: Self-Supervised Self-Supervision (S4): S4 builds on deep probabilistic logic and uses probabilistic logic to represent self-supervision for learning a deep neural network for the end prediction task. Starting from pre-specified self-supervision, S4 interleaves structure learning and active learning steps to introduce new self-supervision for training the neural network and refining the graphical model parameters. Self-supervision factors from initialization, structure learning, and active learning are shown in blue, red, and black, respectively.

only, and used them to simulate oracle self-supervision for initial self-supervision and active learning. We find that S4 can substantially improve over the seed self-supervision by proposing new virtual evidence, and can match the accuracy of fully supervised systems with a fraction of human effort.

Deep Probabilistic Logic

Given a prediction task, let \mathcal{X} and \mathcal{Y} denote the sets of possible inputs and outputs, respectively. The goal is to train a prediction module $\Psi(x, y)$ that scores output y given input x . We assume that $\Psi(x, y)$ represents the conditional probability $P(y|x)$. Let $X = (X_1, \dots, X_N)$ denote a sequence of inputs and $Y = (Y_1, \dots, Y_N)$ the corresponding outputs. If Y is observed, $\Psi(x, y)$ can be learned using standard supervised learning. In this paper, we consider the setting where Y is unobserved, and $\Psi(x, y)$ is learned using self-supervision.

The key idea of deep probabilistic logic (DPL) is to represent self-supervision as prior belief over the latent label variables Y and their interdependencies by combining probabilistic logic and deep learning (Wang and Poon 2018). Pearl (1988) first introduced *virtual evidence* to represent prior belief on the value of a random variable. Specifically, the prior belief on Y can be represented by introducing a binary variable v as a dependent of Y such that $P(v = 1|Y = y)$ is proportional to the belief of $Y = y$. The virtual evidence $v = 1$ can be viewed as a reified variable representing the unary potential function $\Phi(y) \propto P(v = 1|y)$. More generally, this can represent arbitrary potential functions $\Phi(X, Y)$ over the inputs and outputs, so as to model prior beliefs over arbitrary high-order factors. DPL uses Markov logic (Richardson and Domingos 2006) to represent virtual evidences and uses a deep neural network as the prediction module $\Psi(x, y)$. Let $\mathbf{V} = \{v_1, \dots, v_k\}$ be a set of pre-specified virtual evidences, with the corresponding potential functions being (Φ_1, \dots, Φ_k) . We will use K as shorthand

for the event $\mathbf{V} = \mathbf{1}$ (i.e., $v_1 = 1, \dots, v_k = 1$). DPL defines a probability distribution over K, X, Y by combining a factor graph with potentials Φ representing $P(K|Y, X)$ and the prediction module $P(Y|X)$:

$$P(K, Y|X) \propto \prod_v \Phi_v(X, Y) \cdot \prod_i \Psi(X_i, Y_i) \quad (1)$$

Here, the potential functions $\Phi_v(X, Y)$ are represented by weighted first-order logic formulas (i.e., $\Phi_v(X, Y) = \exp(w_v f_v(X, Y))$, with $f_v(X, Y)$ being a binary feature represented by a first-order logical formula). DPL considers a Bayesian setting where each w_v is drawn from a pre-specified prior distribution $w_v \sim P(w_v|\alpha_v)$. Fixed w_v amounts to the special case when the prior is concentrated on the preset value. For uncertain w_v 's, DPL computes their maximum a posteriori (MAP) estimates.

Parameter learning in DPL maximizes the conditional log likelihood of virtual evidences $\log P(K|X)$, which can be done using variational EM. In the E-step, DPL computes a variational approximation $q(Y)$ for $P(Y|K, X)$ using loopy belief propagation (Murphy, Weiss, and Jordan 1999) with current parameters Φ, Ψ , by conducting message passing in $P(K, Y|X)$ iteratively. In the M-step, DPL treats $q(Y)$ as the probabilistic label distribution to train Φ and Ψ via standard supervised learning. For the prediction module Ψ , this reduces to standard deep learning, with the marginals $q_i(Y_i) = \mathbb{E}_{q(Y)}(Y_i)$ serving as probabilistic labels for X_i . For the supervision module, this reduces to standard parameter learning for log-linear models (i.e., learning non-fixed w_v 's), and can be solved using gradient descent, with the partial derivative for w_v being $\mathbb{E}_{\Phi(Y, X)} [f_v(X, Y)] - \mathbb{E}_{q(Y)} [f_v(X, Y)]$. The second expectation can be done by simple counting. The first expectation, on the other hand, requires probabilistic inference in the graphical model. But it can be computed using belief propagation, similar to the E-

Algorithm 1 Self-Supervised Self-Supervision (S4)

Input: Seed virtual evidences I , deep neural network Ψ , inputs $X = (X_1, \dots, X_N)$, unobserved outputs $Y = (Y_1, \dots, Y_N)$, human query budget T .
Output: Learned prediction module Ψ and virtual evidences $K = \{(f_v(X, Y), w_v) : v\}$.
Initialize: $K = I$; $Q = \emptyset$; $i = 0$.
for $i = 1 \dots M$ **do**
 while Structured Self-Training not converged **do**
 $\Psi, K \leftarrow \text{DPL-Learn}(K, \Psi, X, Y)$
 $v = \text{PropSST}(K, \Psi, X, Y)$;
 $K \leftarrow K \cup v$;
 end while
 if $|Q| < T$ **then**
 $v = \text{PropFAL}(K, \Psi, X, Y, Q)$; $Q \leftarrow Q \cup v$;
 if Human-Accept(v) **then** $K \leftarrow K \cup v$
 end if
end for

step, except that the messages are limited to factors in the supervision module (i.e., messages from Ψ are not included)¹.

Self-Supervised Self-Supervision

In this section, we present the *Self-Supervised Self-Supervision (S4)* framework, which extends deep probabilistic logic (DPL) with the capability to learn new self-supervision. Let $\mathcal{V} = \{(f_v, w_v, \alpha_v) : v\}$ be the set of all candidate virtual evidences, where $f_v(X, Y)$ is a first-order logical formula, w_v is the weight, and α_v is the weight prior (for non-fixed w_v). Let K be the set of virtual evidences maintained by the algorithm, initialized by the seed I . The key idea of S4 is to interleave structure learning and active learning to iteratively propose new virtual evidence $v \in \mathcal{V}$ to augment K (Figure 1). Self-training is a special case where candidate virtual evidences are individual label assignments (i.e., $f_v = I[y_v = l_v]$). S4 can thus be viewed as conducting *structured self-training (SST)* by generalizing self-training to admit arbitrary Markov logic formulas as virtual evidence.

S4 can also be viewed as conducting *structure learning* in the factor graph that specifies the virtual evidence. Structure learning has been studied intensively in the graphical model literature (Friedman and Koller 2003). It is also known as feature selection or feature induction in general machine learning literature (Hall 1999). Here, we are introducing structured factors for self-supervision, rather than as feature templates to be used during training. Another key difference from standard structure learning is the deep neural network, which provides an alternative view from the virtual evidence space and enables multi-view learning in DPL. The neural network can also help identify candidate virtual evidences, e.g., via neural attention.

In data programming and many other prior methods, human experts need to pre-specify all self-supervision upfront.

¹In theory, to optimize a lower bound on the conditional log likelihood as in standard EM, E-steps and M-steps need to use different approximate inference algorithms (Domke 2013). However, using the same approximation for both has been shown to work well in practice (Verbeek and Triggs 2008).

While it is easy to generate a small seed by identifying the most salient self-supervision, this effort can quickly become tedious and more challenging as the experts are required to enumerate the less salient templates. On the other hand, given a candidate, it’s generally much easier for experts to validate it. *This suggests that for the best utilization of human bandwidth, we should focus on leveraging them to produce the initial self-supervision and verify candidate self-supervision.* Consequently, in addition to structured self-training (SST), S4 incorporates *feature-based active learning (FAL)* (i.e., active learning of self-supervision). When SST converges, S4 will switch to the active learning mode by proposing a candidate virtual evidence for human verification (i.e., labeling a feature rather than an instance in standard active learning). Intuitively, in FAL we are proposing virtual evidences for which the labels of the corresponding instances are still uncertain. If the human expert can provide definitive supervision on the label, the information gain will be large. By contrast, in SST, we favor virtual evidences with skewed posterior label assignments for their corresponding instances, as they can potentially amplify the signal.

Algorithm 1 describes the S4 algorithm. S4 first conducts DPL using the initial self-supervision I , then interleaves structured self-training (SST) with feature-based active learning (FAL). SST steps are repeated until there is little change in the probabilistic labels (less than 1% in our experiments). DPL learning updates the deep neural network and the graphical model parameters with warm start (i.e., the parameters are initialized with the previous parameters). All proposed queries are stored and won’t be proposed again. The total amount of human effort consists of generating the seed I and validating T queries in active learning.

S4 is a general algorithm framework that can combine various strategies for designing \mathcal{V} , PropSST, and PropFAL. In standard structure learning, PropSST would attempt to maximize the learning objective (e.g., conditional likelihood of seed virtual evidences) by iteratively conducting greedy structure changes. However, this is very expensive to compute, since it requires a full DPL run just to score each candidate. Instead, we take inspiration from the feature-induction and relational learning literature and use heuristic approximations that are much faster to evaluate. In the most general setting, \mathcal{V} contains all possible potential functions. In practice, we can restrict it to a tractable subset to obtain a good trade-off between expressiveness and computation for the problem domain. Interestingly, as we will see in the experiment section, even with relatively simple classes of self-supervision, S4 can dramatically improve over DPL through structure learning and active learning.

We use text classification from natural language processing (NLP) as a running example to illustrate how to apply S4 in practice. Here, the input $X_i = (t_1, \dots, t_{s_i})$ is a sequence of tokens and the output Y_i is the classification label (e.g., pos or neg in sentiment analysis).

Candidate Self-Supervision

For \mathcal{V} , the simplest choice is to use tokens. Namely, $f_{t,l}(X_i, Y_i) = \mathbb{I}[t \in X_i \wedge Y_i = l]$. For simplicity, we can use a fixed weight and prior for all initial virtual evidence, i.e.,

$\mathcal{V} = \{(f_{t,l}, w, \alpha) : t, l\}$. Take sentiment analysis as an example. X_i may represent a movie review and $Y_i \in \{0, 1\}$ the sentiment. A virtual evidence for self-supervision may stipulate that if the review contains the word “good”, the sentiment is more likely to be positive. This can be represented by the formula $f_{\text{good},1}(X_i, Y_i) = \mathbb{I}[\text{“good”} \in X_i \wedge Y_i = 1]$ with a positive weight. A more advanced choice for \mathcal{V} may include high-order factors, such as $f_{ij}(Y_i, Y_j) = \mathbb{I}[Y_i = Y_j]$. If we add this factor for similar pairs X_i, X_j , it stipulates that instances with similar input are likely to share the same label. Here we define similar pairs with a similarity function $\text{Sim}(X_i, X_j)$ between X_i and X_j , such as the cosine similarity between the sentence (or document) embeddings of X_i and X_j , based on the current deep neural network. Note that this is different from graph-based semi-supervised learning or other kernel-based methods in that the similarity metric is not pre-specified and fixed, but rather evolving along with the deep neural network for the end task.

Structured Self-Training (PropSST)

From DPL learning, we obtain the current marginal estimate of the latent label variables $q_i(Y_i)$, which we would treat as probabilistic labels in assessing candidate virtual evidence. There are many sensible strategies for proposing candidates in structured self-training (i.e., PropSST). For token-based self-supervision, a common technique from the feature-selection literature is to choose a token highly correlated with a label. For example, we can choose the token t that occurs much more frequently in instances for a given label l than others using our noisy label estimates. We find that this often leads to very noisy proposals and semantic drift. A simple refinement is to restrict our scoring to instances containing some initial self-supervised tokens. However, this still has the drawback that a word may occur more often in instances of a class for reasons other than contributing to the label classification. We therefore consider a more sophisticated strategy based on neural attention. Namely, we will credit occurrences using the normalized attention weight for the given token in each instance.

Formally, let $A_\Psi(X_i, j)$ represent the normalized attention weight the neural network Ψ assigns to the j -th token in X_i for the final classification. We define *average weighted attention* for token t and label l as $\text{Att}_n(t, l) = \frac{1}{C_t} \sum_{i,j: X_{i,j}=t} q_i(Y_i = l) \cdot A_\Psi(X_i, j)$, where C_t is the number of occurrences of t in X . Then PropSST would simply score token-based self-supervision $f_{t,l}$ using relative average weighted attention: $S_{\text{token}}(t, l) = \text{Att}_n(t, l) - \sum_{l' \neq l} \text{Att}_n(t, l')$. In each iteration, PropSST picks the top scoring $f_{t,l}$ that has not been proposed yet as the new virtual evidence to add to K .

We also consider an entropy-based score function that works for arbitrary input-based features. It treats the prediction module Ψ as a black box, and only uses the posterior label assignments $q_i(Y_i)$. Consider candidate virtual evidence $f_{b,l}(X_i, Y_i) = \mathbb{I}[b(X_i) \wedge Y_i = l]$, where b is a binary function over input X_i . This clearly generalizes token-based virtual evidence. Define $\text{Ent}(b) = H\left(\frac{1}{C_b} \sum_{i:b(X_i)=1} q_i(Y_i)\right)$, where H is the Shannon entropy and C_b is the number of

instances for which the feature b holds true. This function represents the entropy of the average posterior among all instances with $b(X_i) = 1$. PropSST will then use $S_{\text{entropy}}(b) = 1/\text{Ent}(b)$ to choose the b^* with the lowest average entropy and then pick label l^* with the highest average posterior probability for b^* . In our experiments, this performs similarly to attention-based scores.

For joint-inference self-supervision, we consider the similarity-based factors defined earlier, and leave the exploration of more complex factors to future work. To distinguish task-specific similarity from pretrained similarity, we use the difference between the similarity computed using the current fine-tuned BERT model and that using the pretrained one. Formally, let $\text{Sim}_{\text{pretrained}}(X_i, X_j)$ be the cosine similarity between the embeddings of X_i and X_j generated by the pretrained BERT model, and $\text{Sim}_\Psi(X_i, X_j)$ be that between the embeddings using the current learned network Ψ . PropSST would score the joint-inference factor using the relative similarity and choose the top scoring pairs to add to self-supervision: $S_{\text{joint}}(X_i, X_j) = \text{Sim}_\Psi(X_i, X_j) - \text{Sim}_{\text{pretrained}}(X_i, X_j)$.

Feature-Based Active Learning (PropFAL)

For active learning, a common strategy is to pick the instance with highest entropy in the label distribution based on the current marginal estimate. In feature-based active learning, we can similarly pick the feature b with the highest average entropy $\text{Ent}(b)$. Note that this is opposite to how we use the entropy-based score function in PropSST, where we choose the feature with the lowest average entropy. In PropFAL, we will identify $b^* = \arg \max(\text{Ent}(b))$, present $f_{b^*,l}(X, Y) = \mathbb{I}[b^*(X) \wedge Y = l]$ for all possible labels l , and ask the human expert to choose a label l^* to accept or reject them all.

Experiments

We use the natural language processing (NLP) task of text classification to explore the potential for S4 to improve over DPL using structure learning and active learning. We used three standard text classification datasets: IMDb (Maas et al. 2011), Stanford Sentiment Treebank (Socher et al. 2013), and Yahoo! Answers (Zhang, Zhao, and LeCun 2015). IMDb contains movie reviews with polarity labels (positive/negative). There are 25,000 training instances with equal numbers of positive and negative labels, and the same numbers for test. Stanford Sentiment Treebank (StanSent) also contains movie reviews, but was annotated with five labels ranging from very negative to very positive. We used the binarized version of StanSent, which collapses the polarized categories and discards the neutral sentences. It contains 6,920 training instances and 1,821 test instances, with roughly equal split. Overall, the StanSent reviews are shorter than IMDb’s, and they often exhibit more challenging linguistic phenomena (e.g., nested negations or sarcasm). The Yahoo dataset contains 1.4 million training questions and 60,000 test questions from Yahoo! Answers; these are equally split into 10 classes. Due to space constraints, we report the Yahoo results in the appendix.

In all our experiments with S4, we withheld gold labels from the system, used the training instances as unlabeled

data, and evaluated on the test set. We reported test accuracy, as all of the datasets are class-balanced. For our neural network prediction module $\Psi(X_i, Y_i)$, we used the standard BERT-base model pretrained using Wikipedia (Devlin et al. 2018), along with a global-context attention layer as in Yang et al. (2016), which we also used for attention-based scoring. We truncated the input text to 512 tokens, the maximum allowed by the standard BERT model. All of our baselines (except supervised bag-of-words) use the same BERT model. For all virtual evidences, we used initial weight $w = 2.2$ (the log-odds of 90% probability) and used an α corresponding to an L2 penalty of 5×10^{-8} on w . Our results are not sensitive to these values. In all experiments, we use the Adam optimizer with an initial learning rate tuned over [0.1, 0.01, 0.001]. The optimizer’s history is reset after each EM iteration to remove old gradient information. We always performed 3 EM iterations and trained Ψ for 5 epochs per iteration.

For virtual evidence, we focus on token-based unary factors and similarity-based joint factors, as discussed in the previous section, and leave the exploration of more complex factors to future work. Even with these factors, our self-supervised Ψ models often nearly match the accuracy of the best supervised models. We also compare with Snorkel, a popular self-supervision system (Ratner et al. 2016). We use the latest Snorkel version (Ratner et al. 2019), which models correlations among same-instance factors. Snorkel cannot incorporate joint-inference factors across different instances.

To simulate human supervision for unary factors, we trained a unigram model using the training data with L1 regularization and selected the 100 tokens with the highest weights for each class as the oracle self-supervision. By default, we used the top tokens for each class in the initial self-supervision I . We also experimented with using random tokens from the oracle in I to simulate lower-quality initial supervision and to quantify the variance of S4. For the set of oracle joint factors, we fine-tuned the standard BERT model on the training set, used the CLS embedding BERT produces to compute input similarity, and picked the 100 input pairs whose similarity changed the most between the fine-tuned model and the initial model.

We first investigate whether structure learning can help in S4 by running without feature-based active learning. We set the query budget $T = 0$ in Algorithm 1. Because we only take structured self-training steps when $T = 0$, we denote this version of S4 as S4-SST. Table 1 shows the results on IMDb. With just six self-supervised tokens (three per class), S4-SST already attained 86% test accuracy, which outperforms self-training with 100 labeled examples by 16 absolute points, and is only slightly worse than self-training with 1000 labeled examples or supervised training with 25,000 labeled examples. By conducting structure learning, S4-SST substantially outperformed DPL, gaining about 5 absolute points in accuracy (a 25% relative reduction in error), and also outperformed the Snorkel baseline by 8.9 points. Interestingly, with more self-supervision at 20 tokens, DPL’s performance drops slightly, which might stem from more noise in the initial self-supervision. By contrast, S4-SST capitalized on the larger seed self-supervision and attained steady

Algorithm	Sup. size $ I $	Test acc (%)
BoW	25k	87.1
DNN	25k	91.0
Self-training	100	69.9
Self-training	1k	88.5
Snorkel	6	76.6
DPL	6	80.7
S4-SST	6	85.5
S4 ($T = 20$)	6	85.6
Snorkel	20	82.4
DPL	20	78.9
S4-SST	20	86.4
S4 ($T = 20$)	20	86.9
Snorkel	40	84.6
DPL	40	85.2
S4-SST	40	86.6
S4 ($T = 20$)	40	86.8

Table 1: System comparison on IMDb

improvement. Even with substantially more self-supervision at 40 tokens, S4-SST still attained similar accuracy gain, demonstrating the power in structure learning. On average across different initial amounts of supervision $|I|$, S4-SST outperforms DPL by 5.6 points and Snorkel by 5.3 points. Next, we consider the full S4 algorithm with a budget of up to $T = 20$ human queries (S4 ($T = 20$)). Overall, by automatically generating self-supervision from structure learning, S4-SST already attained very high accuracy on this dataset. However, active learning can still produce some additional gain. The only randomness in Table 1 is the initialization of the deep network Ψ , which has a negligible effect.

Figure 2 (leftmost) shows how S4-SST iterations improve the test accuracy of the learned neural network with different amounts of initial virtual evidence. Not surprisingly, with more initial self-supervision, the gain is less pronounced, but still significant. Figure 2 (center left) compares the learning curves of S4-SST with those of self-training. Remarkably, with just *six* self-supervised tokens, S4-SST not only attained substantial gain over the iterations, but also easily outperformed self-training despite the latter using an order of magnitude more label information (up to 200 labeled examples). This shows that S4 is much more effective in leveraging bounded human effort for supervision.

Figure 2 (center right) shows 10 runs of S4-SST when it was initialized with 20 *random* oracle tokens (10 per class), rather than the *top* 20 tokens from the oracle. As expected, DPL’s initial performance is worse than with the top oracle tokens. However, over the iterations, S4-SST was able to recover even from particularly poor initial state, gaining up to 20 absolute accuracy points over DPL in the process. The final accuracy of S4-SST is 85.2 ± 0.9 , compared to 71.5 ± 6.5 for DPL and 72.9 ± 6.7 for Snorkel, a mean improvement of more than 12 absolute accuracy points over both baselines. S4-SST’s gains over DPL and Snorkel are statistically significant using a paired *t*-test (samples are paired

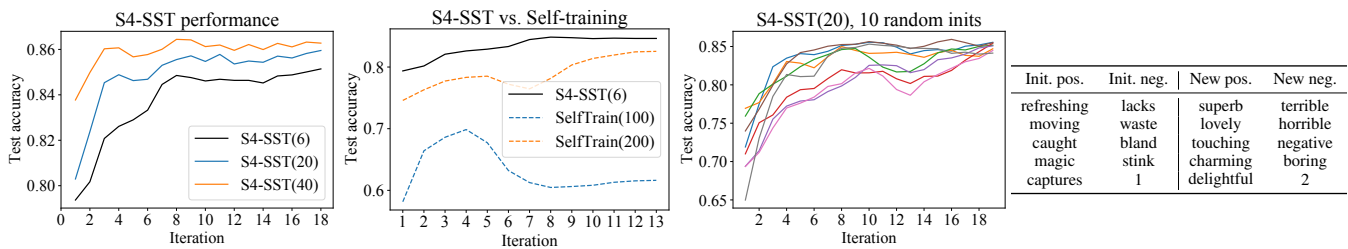


Figure 2: S4-SST learning in IMDb. Left-to-right: S4-SST learning curves with various numbers of initial self-supervised tokens; comparison between the learning curves of S4-SST with just six self-supervised tokens and self-training with an order of magnitude more direct supervision (100/200 labeled examples); multiple runs of S4-SST with 20 *random* oracle self-supervised tokens (rather than the top oracle tokens); pre-specified self-supervised tokens and proposed tokens in the first few iterations.

Algorithm	Sup. size $ I $	Test acc (%)
BoW	6.9k	78.9
DNN	6.9k	90.9
Self-training	50	78.0
	100	81.8
Snorkel	6	63.5
DPL	6	57.2
S4-SST	6	73.0
S4-SST + J	6	76.2
S4 + J ($T = 20$)	6	81.4
Snorkel	20	73.0
DPL	20	72.4
S4-SST	20	83.3
S4-SST + J	20	85.1
S4 + J ($T = 20$)	20	84.4
Snorkel	40	73.6
DPL	40	77.0
S4-SST	40	84.9
S4-SST + J	40	86.3
S4 + J ($T = 20$)	40	85.4

Table 2: System comparison on Stanford

when the algorithms have the same initial factors I) with $p = 0.01$. This indicates that S4-SST is robust to noise in the initial self-supervision. Figure 2 (rightmost) shows the initial self-supervised tokens and the ones learned by S4-SST in the first few iterations. We can see that S4-SST is able to learn highly relevant self-supervision tokens.

S4 has similar gains over DPL, Snorkel, and self-training on the Stanford dataset. See Table 2. The Stanford dataset is much more challenging and for a long while, it was hard to exceed 80% test accuracy (Socher et al. 2013). Interestingly, S4-SST was able to surpass this milestone using just 20 initial self-supervision tokens. As in Table 1, the only randomness is in the initialization of Ψ , which is negligible.

For IMDb, as can be seen above, even with simple token-based self-supervision, S4-SST already performed extremely well. So we focused our investigation of joint-inference factors on the more challenging Stanford dataset (S4-SST + J). While S4-SST already performed very well

	$ I = 6$	$ I = 10$	$ I = 20$	$ I = 40$
$T = 5$	79.0	83.8	84.0	85.9
$T = 10$	79.5	83.7	85.0	86.0
$T = 20$	82.7	84.5	85.8	86.4

Table 3: S4-FAL results on IMDb (no SST steps)

	$ I = 6$	$ I = 10$	$ I = 20$	$ I = 40$
$T = 5$	71.4	77.5	82.5	83.9
$T = 10$	72.3	77.4	82.5	83.7
$T = 20$	77.1	80.8	83.4	84.0

Table 4: S4-FAL results on Stanford (no SST steps)

with token-based self-supervision, by incorporating joint-inference factors, it could still gain up to 3 absolute accuracy points. An example learned joint-inference factor is between the sentence pair: *This is no "Waterboy!"* and *"It manages to accomplish what few sequels can—it equals the original and in some ways even better"*. Note that Waterboy is widely considered a bad movie, hence the first sentence expresses a positive sentiment, just like the second. It is remarkable that S4 can automatically induce such factors with complex and subtle semantics from small initial self-supervision.

The Stanford results also demonstrated that active learning could play a bigger role in more challenging scenarios. With limited initial self-supervision ($|I| = 6$), the full S4 system (S4+J ($T=20$)) gained 8 absolute points over S4-SST and 5 absolute points over S4-SST+J. With sufficient initial self-supervision and joint-inference, however, active learning was actually slightly detrimental ($|I| = 20, 40$).

Finally, we evaluate S4-FAL, which conducts active learning but not structure learning. See Tables 3 and 4. As expected, performance improved with larger initial self-supervision (I) and human query budget (T). Active learning helps the most when initial self-supervision is limited. Compared to S4 with structure learning, however, active learning alone is less effective. For example, without requiring any human queries, S4-SST outperformed S4-FAL on both IMDB and Stanford even when the latter was allowed up to $T = 20$ human queries.

Related Work

Techniques to compensate for the lack of direct supervision come in many names and forms (Mintz et al. 2009; Ratner et al. 2016; Bach et al. 2017; Roth 2017; Wang and Poon 2018). Self-supervision has emerged as an encompassing paradigm that views these as instances of using self-specified templates to generate noisy labeled examples on unlabeled data. The name *self-supervision* is closely related to *self-training* (McClosky, Charniak, and Johnson 2006), which bootstraps from a supervised classifier, uses it to annotate unlabeled instances, and iteratively uses the confident labels to retrain the classifier. *Task-agnostic self-supervision* generalizes word embedding and language modeling by learning to predict self-specified masked tokens, as exemplified by recent pretraining methods such as BERT (Devlin et al. 2018). In this paper, we focus on *task-specific self-supervision* and use pretrained models as a building block for task-specific learning.

Existing self-supervision paradigms are typically special cases of deep probabilistic logic (DPL). E.g., the popular *data programming* methods (Ratner et al. 2016; Bach et al. 2017; Varma et al. 2017) admit only virtual evidences for individual instances (labeling functions or their correlations). *Anchor learning* (Halpern et al. 2016) is an earlier form of data programming that, while more restricted, allows for stronger theoretical learning guarantees. *Prototype learning* is an even earlier special case with labeling functions provided by “prototypes” (Haghighi and Klein 2006; Poon 2013). Using Markov logic to model self-supervision, DPL can incorporate arbitrary prior beliefs on both individual labels and their interdependencies, thereby unleashing the full power of *joint inference* (Chang, Ratinov, and Roth 2007; Druck, Mann, and McCallum 2008; Poon and Domingos 2008; Ganchev et al. 2010) to amplify and propagate self-supervision signals.

Self-supervised self-supervision (S4) further extends DPL with structure learning capability. Most structure learning techniques are developed for the supervised setting, where structure search is guided by labeled examples (Friedman and Koller 2003; Kok and Domingos 2005). Moreover, traditional relational learning induces deterministic rules and is susceptible to noise and uncertainty. *Bootstrapping* learning is one of the earliest and simplest self-supervision methods with some rule-learning capability, by alternating between inducing characteristic contextual patterns and classifying instances (Hearst 1992; Carlson et al. 2010). The pattern classes are limited and only applicable to special problems (e.g., “A such as B” to find ISA relations). Most importantly, they lack a coherent probabilistic formulation and may suffer catastrophic *semantic drift* due to ambiguous patterns (e.g., “cookie” as food or compute use). Yarowsky (1995) and Collins and Singer (1999) designed a more sophisticated rule induction approach, but their method uses deterministic rules and may be sensitive to noise and ambiguity. Recently, Snuba (Varma and Ré 2018) extends the data programming framework by automatically adding new labeling functions, but like prior data programming methods, their self-supervision framework is limited to modeling prior beliefs on individual instances. Their method also re-

quires access to a small number of labeled examples to score new labeling functions.

Another significant advance in S4 is by extending DPL with the capability to conduct *structured active learning*, where human experts are asked to verify arbitrary virtual evidences, rather than a label decision. Note that by admitting joint inference factors, this is more general than prior use of *feature-based active learning*, which focuses on per-instance features (Druck, Settles, and McCallum 2009). As our experiments show, interleaving structured self-training learning and structured active learning results in substantial gains, and provides the best use of precious human bandwidth. Tong and Koller (2001) previously considered active structure learning in the context of Bayesian networks. Anchor learning (Halpern et al. 2016) can also suggest new self-supervision for human review. Darwin (Galhotra, Golshan, and Tan 2020) incorporates active learning for verifying proposed rules, but it doesn’t conduct structure learning, and like Snuba and other data programming methods, it only models individual instances.

Neural-symbolic learning and reasoning has received increasing attention (Besold et al. 2017). In particular, combining deep learning with probabilistic models can leverage their complementary strengths in modeling complex patterns and infusing rich prior knowledge. Prior work tends to focus on deep generative models that aim to uncover latent factors for generative modeling and semi-supervised learning (Kingma and Welling 2013; Kingma et al. 2014). They admit limited forms of self-supervision (e.g., latent structures such as Markov chains (Johnson et al. 2016)). S4 and DPL instead combine a discriminative neural network predictor with a generative self-supervision model based on Markov logic, and can fully leverage their respective capabilities to advance co-learning (Blum and Mitchell 1998; Grechkin, Poon, and Howe 2017). Deep neural networks also provide a powerful feature-induction engine to support structure learning and active learning.

Conclusion

We present Self-Supervised Self-Supervision (S4), a general self-supervision framework that can automatically induce new self-supervision by extending deep probabilistic logic (DPL) with structure learning and active learning capabilities. Our experiments on various natural language processing (NLP) tasks show that compared to prior systems for task-specific self-supervision, such as Snorkel and DPL, S4 can obtain gain up to 20 absolute accuracy points with the same amount of supervision. S4 only relies on humans to identify the most salient self-supervision for initialization and to verify proposed self-supervision, which tends to be the most effective use of human bandwidth. While we focus on NLP tasks in this paper, our methods are general and can potentially be applied to other domains. Future directions include: further investigation in combining structure learning and active learning; exploring more sophisticated self-supervision classes and proposal algorithms; applications to other domains.

Acknowledgments

We give warm thanks to Naoto Usuyama, Cliff Wong, and the rest of Project Hanover team for helpful discussions.

References

- Bach, S. H.; He, B.; Ratner, A.; and Ré, C. 2017. Learning the structure of generative models without labeled data. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 273–282.
- Besold, T. R.; d’Avila Garcez, A.; Bader, S.; Bowman, H.; Domingos, P.; Hitzler, P.; Kühnberger, K.-U.; Lamb, L. C.; Lowd, D.; Lima, P. M. V.; de Penning, L.; Pinkas, G.; Poon, H.; and Zaverucha, G. 2017. Neural-symbolic learning and reasoning: A survey and interpretation. In *arxiv*.
- Blum, A.; and Mitchell, T. 1998. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, 92–100.
- Carlson, A.; Betteridge, J.; Kisiel, B.; Settles, B.; Hruschka, E. R.; and Mitchell, T. M. 2010. Toward an architecture for never-ending language learning. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- Chang, M.-W.; Ratinov, L.; and Roth, D. 2007. Guiding semi-supervision with constraint-driven learning. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, 280–287.
- Collins, M.; and Singer, Y. 1999. Unsupervised models for named entity classification. In *1999 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Domke, J. 2013. Learning graphical model parameters with approximate marginal inference. *IEEE transactions on pattern analysis and machine intelligence* 35(10): 2454–2467.
- Druck, G.; Mann, G.; and McCallum, A. 2008. Learning from labeled features using generalized expectation criteria. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, 595–602.
- Druck, G.; Settles, B.; and McCallum, A. 2009. Active learning by labeling features. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*, 81–90. Association for Computational Linguistics.
- Friedman, N.; and Koller, D. 2003. Being Bayesian about network structure. A Bayesian approach to structure discovery in Bayesian networks. *Machine learning* 50(1-2): 95–125.
- Galhotra, S.; Golshan, B.; and Tan, W.-C. 2020. Adaptive Rule Discovery for Labeling Text Data. *arXiv preprint arXiv:2005.06133*.
- Ganchev, K.; Graça, J.; Gillenwater, J.; and Taskar, B. 2010. Posterior regularization for structured latent variable models. *Journal of Machine Learning Research* 11(67): 2001–2049.
- Grechkin, M.; Poon, H.; and Howe, B. 2017. Ezlearn: Exploiting organic supervision in large-scale data annotation. *arXiv preprint arXiv:1709.08600*.
- Haghighi, A.; and Klein, D. 2006. Prototype-driven learning for sequence models. In *Proceedings of the main conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics*, 320–327. Association for Computational Linguistics.
- Hall, M. A. 1999. Correlation-based feature selection for machine learning.
- Halpern, Y.; Horng, S.; Choi, Y.; and Sontag, D. 2016. Electronic medical record phenotyping using the anchor and learn framework. *Journal of the American Medical Informatics Association* 23(4): 731–740.
- Hearst, M. A. 1992. Automatic acquisition of hyponyms from large text corpora. In *Proceedings of the 14th conference on Computational linguistics-Volume 2*, 539–545. Association for Computational Linguistics.
- Johnson, M. J.; Duvenaud, D. K.; Wiltchko, A.; Adams, R. P.; and Datta, S. R. 2016. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in neural information processing systems*, 2946–2954.
- Kingma, D. P.; Mohamed, S.; Rezende, D. J.; and Welling, M. 2014. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, 3581–3589.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kok, S.; and Domingos, P. 2005. Learning the structure of Markov logic networks. In *Proceedings of the 22nd international conference on Machine learning*, 441–448.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature* 521(7553): 436–444.
- Maas, A. L.; Daly, R. E.; Pham, P. T.; Huang, D.; Ng, A. Y.; and Potts, C. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1*, 142–150. Association for Computational Linguistics.
- McClosky, D.; Charniak, E.; and Johnson, M. 2006. Effective self-training for parsing. In *Proceedings of the main conference on human language technology conference of the North American Chapter of the Association of Computational Linguistics*, 152–159. Association for Computational Linguistics.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual*

- Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2, 1003–1011. Association for Computational Linguistics.
- Murphy, K. P.; Weiss, Y.; and Jordan, M. I. 1999. Loopy belief propagation for approximate inference: an empirical study. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, 467–475.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.
- Poon, H. 2013. Grounded unsupervised semantic parsing. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 933–943.
- Poon, H.; and Domingos, P. 2008. Joint unsupervised coreference resolution with Markov logic. In *Proceedings of the 2008 conference on empirical methods in natural language processing*, 650–659.
- Ratner, A.; Hancock, B.; Dunnmon, J.; Sala, F.; Pandey, S.; and Ré, C. 2019. Training complex models with multi-task weak supervision. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 4763–4771.
- Ratner, A. J.; De Sa, C. M.; Wu, S.; Selsam, D.; and Ré, C. 2016. Data programming: Creating large training sets, quickly. In *Advances in neural information processing systems*, 3567–3575.
- Richardson, M.; and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1-2): 107–136.
- Roth, D. 2017. Incidental supervision: Moving beyond supervised learning. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A. Y.; and Potts, C. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, 1631–1642.
- Tong, S.; and Koller, D. 2001. Active learning for structure in Bayesian networks. In *International joint conference on artificial intelligence*, volume 17, 863–869. Citeseer.
- Varma, P.; He, B. D.; Bajaj, P.; Khandwala, N.; Banerjee, I.; Rubin, D.; and Ré, C. 2017. Inferring generative model structure with static analysis. In *Advances in neural information processing systems*, 240–250.
- Varma, P.; and Ré, C. 2018. Snuba: automating weak supervision to label training data. *Proceedings of the VLDB Endowment* 12(3): 223–236.
- Verbeek, J.; and Triggs, B. 2008. Scene segmentation with conditional random fields learned from partially labeled images. In *Proc. NIPS*.
- Wang, H.; and Poon, H. 2018. Deep Probabilistic Logic: A Unifying Framework for Indirect Supervision. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 1891–1902.
- Yang, Z.; Yang, D.; Dyer, C.; He, X.; Smola, A.; and Hovy, E. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, 1480–1489.
- Yarowsky, D. 1995. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics*, 189–196.
- Zhang, X.; Zhao, J.; and LeCun, Y. 2015. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, 649–657.