



Published in final edited form as:

Adv Neural Inf Process Syst. 2017 December ; 30: 239–249.

Inferring Generative Model Structure with Static Analysis

Paroma Varma, Bryan He, Payal Bajaj, Imon Banerjee, Nishith Khandwala, Daniel L. Rubin, and Christopher Ré
Stanford University

Abstract

Obtaining enough labeled data to robustly train complex discriminative models is a major bottleneck in the machine learning pipeline. A popular solution is combining multiple sources of weak supervision using generative models. The structure of these models affects training label quality, but is difficult to learn without any ground truth labels. We instead rely on these weak supervision sources having some structure by virtue of being encoded programmatically. We present Coral, a paradigm that infers generative model structure by statically analyzing the code for these heuristics, thus reducing the data required to learn structure significantly. We prove that Coral's sample complexity scales quasilinearly with the number of heuristics and number of relations found, improving over the standard sample complexity, which is exponential in n for identifying n^{th} degree relations. Experimentally, Coral matches or outperforms traditional structure learning approaches by up to 3.81 F1 points. Using Coral to model dependencies instead of assuming independence results in better performance than a fully supervised model by 3.07 accuracy points when heuristics are used to label radiology data without ground truth labels.

1 Introduction

Deep neural networks and other complex discriminative models rely on a large amount of labeled training data for their success. For many real-world applications, obtaining this magnitude of *labeled* data is one of the most expensive and time consuming aspects of the machine learning pipeline. Recently, generative models have been used to create training labels from various weak supervision sources, such as heuristics or knowledge bases, by modeling the true class label as a latent variable. After the necessary parameters for the generative models are learned using unlabeled data, the distribution over the true labels can be inferred. Properly specifying the structure of these generative models is essential in estimating the accuracy of the supervision sources. While traditional structure learning approaches have focused on the supervised case [35, 29, 48], previous works related to weak supervision assume that the structure is user-specified [1, 44, 38, 34]. Recently, Bach et al. [2] showed that it is possible to learn the structure of these models with a sample complexity that scales sublinearly with the number of possible binary dependencies. However, the sample complexity scales exponentially for higher degree dependencies, limiting its ability to learn complex dependency structures. Moreover, the time required to learn the dependencies also grows exponentially with the degree of dependencies, hindering the development of user-defined heuristics.

This poses a problem in many domains, where high degree dependencies are common among heuristics that operate over a shared set of inputs. These inputs are interpretable characteristics extracted from the data. For example, various approaches in computer vision use bounding box and segmentation attributes [24, 36, 23], like location and size, to weakly supervise more complex image-based learning tasks [13, 45, 6, 32, 8]. Another example comes from the medical imaging domain, where attributes include characteristics such as the area, intensity and perimeter of a tumor, as shown in Figure 1. Note that these attributes and the heuristics written over them are encoded programmatically. There typically is a relatively small set of interpretable characteristics, so the heuristics often share these attributes. This results in high order dependency structures among these sources, which are crucial to model in the generative model that learns accuracies for these sources.

To efficiently learn higher order dependencies, we present Coral, a paradigm that statically analyzes the source code of the weak supervision sources to infer, rather than learn, the complex relations among heuristics. Coral’s sample complexity scales quasilinearly with the number of relevant dependencies. Moreover, the time to identify these relations is constant in the degree of dependencies, since it only requires looking at the source code for each heuristic. Specifically, Coral analyzes the code used to generate the weak supervision heuristics and its inputs to find which heuristics share the same inputs. This information is used to generate a dependency structure for the heuristics, and a generative model learns the proper weights for this structure to assign probabilistic labels to training data.

We experimentally validate the performance of Coral across various domains and show that it outperforms traditional structure learning under various conditions while being significantly more computationally efficient. We show how modeling dependencies leads to an improvement of 3.81 F1 points compared to standard structure learning approaches. Additionally, we show that Coral can assign labels to data that has no ground truth labels, and this augmented training set results in improving the discriminative model performance by 3.07 points. For a complex relation-based image classification task, 6 heuristic functions written using *only* object label and location as primitives are able to train a model that comes within 0.74 points of the F1 score achieved by a fully-supervised model trained on the rich, hand-labeled attribute and relation information in the Visual Genome database [26].

2 The Coral Paradigm

The Coral paradigm takes as input a set of domain-specific primitives and a set of programmatic user-defined heuristic functions that operate over the primitives. We formally define these abstractions in Section 2.1. Coral runs static analysis on the source code that creates the primitives and the heuristic functions to identify which sets of heuristics are related by virtue of sharing primitives (Section 2.2). Once Coral identifies these dependencies, it uses a factor graph to model the relationship between the heuristics, primitives and the true class label. We describe the conditions under which Coral can learn the structure of the generative model with significantly less data than traditional approaches in Section 2.3 and show how this affects generative model accuracy via simulations. Finally, we discuss how Coral learns the accuracies of each heuristic and outputs probabilistic labels for the training data (Section 2.4).

2.1 Coral Abstractions

Domain-Specific Primitives—Domain-specific primitives (DSPs) in Coral are the simplest elements that heuristic functions take as input and operate over. DSPs in Coral have semantic meaning, making them interpretable for users. This is akin to the concept of language primitives in programming languages, where they are the smallest unit of processing with meaning. The motivation for making the DSPs domain-specific instead of a general construct for the various data modalities is to allow users to take advantage of existing work in their field, which extracts meaningful characteristics from the raw data.

Figure 1 shows an example of a pipeline for bone tumor classification as malignant or benign, inspired by one of our real experiments. First, an automated segmentation algorithm is used to generate a binary mask for the tumor [46, 25, 31, 41]. Then, we extract three DSPs from the segmentation: $\text{area}(p_1)$, $\text{perimeter}(p_2)$ and $\text{total intensity}(p_3)$ of the segmented area. More complex characteristics such as texture, shape and edge features can also be used [18, 4, 27] (see Appendix).

We now define a formal construct for programmatically encoding DSPs. Users generate DSPs in Coral through a primitive specifier function, such as `create_primitives` in Figure 2(a). Specifically, this function takes as input a single unlabeled data point (and necessary intermediate representations such as the segmentation) and returns an instance of `PrimitiveSet`, which maps primitive names to primitive values, like integers (we refer to a specific instance of this class as \mathcal{P}). Note that $\mathcal{P}.\text{ratio}$ is composed of two other primitives, while the rest are generated independently from the image and segmentation.

Heuristic Functions—In Coral, heuristic functions (HFs) can be viewed as a mapping from a subset of the DSPs to a noisy label for the training data, as shown in Figure 1. In our experience with user-defined HFs, we note that they are mostly nested if-then statements where each statement checks whether the value of a single primitive or a combination of them are above or below a user-set threshold (see Appendix). As shown in Figure 3(a), they take as input fields of the object \mathcal{P} and return a label based on the value of the input primitives. While our running example focuses on a single data point for DSP generation and HF construction, both procedures are applied to the entire training set to assign a set of noisy labels from each HF to each data point.

2.2 Static Dependency Analysis

Since the number of DSPs in some domains are relatively small, multiple HFs can operate over the same DSPs, as shown with λ_2 and λ_3 in Figure 1 and Figure 3(a). HFs that share at least one primitive are trivially related to each other. Prior work learns these dependencies using the labels HFs assign to data points and its probability of success scales with the amount of data available [2]. However, it can only learn dependencies among pairs of HFs efficiently, since the data required grows exponentially with the degree of the HF relation. This in turn limits the complexity of the dependency structure this method can accurately learn and model.

Heuristic Function Inputs—Coral takes advantage of the fact that users write HFs over a known, finite set of primitives. It *infers* dependencies that exist among HFs by simply looking at the source code of how the DSPs and HFs are constructed. This process *requires no data* to successfully learn the dependencies, making it more computationally efficient than standard approaches. In order to determine whether any set of HFs share at least one DSP, Coral looks at the input for each HF. Since the HFs only take as input the DSP they operate over, simply grouping HFs by the primitives they share is an efficient approach for recognizing these dependencies. As shown in our running example, this would result in Coral not recognizing any dependencies among the HFs since the input for all three HFs are different (Figure 3(a)). This, however, would be incorrect, since the primitive `P.ratio` is composed of `P.perimeter` and `P.intensity`, which makes λ_2 and λ_3 related. Therefore, along with looking at the primitives that each HF takes as input, it is also essential to model how these primitives are *composed*.

Primitive Compositions—We use our running example in Figure 2 to explain how Coral gathers information about DSP compositions. Coral builds an abstract syntax tree (AST) to represent computation the `create_primitives` function performs. An AST represents operations involving the primitives as a tree, as shown in Figure 2(b).

To find primitive compositions from the AST, Coral first finds the expressions in the AST that add primitives to `P` (denoted in the AST as `P.name`). Then, for each assignment expression, Coral traverses the subtree rooted at the assignment expression and adds all other encountered primitives as a dependency for `P.name`. If no primitives are encountered in the subtree, the primitive is registered as being independent of the rest. The composition structure that results from traversing the AST is shown in Figure 2(c), where `P.area`, `P.intensity`, and `P.perimeter` are independent while `P.ratio` is a composition.

Heuristic Function Dependency Structure—With knowledge of how the DSPs are composed, we return to our original method of looking at the inputs of the HFs. As before, we identify that λ_1 and λ_2 use `P.area` and `P.perimeter`, respectively. However, we now know that λ_3 uses `P.ratio`, which is a composition of `P.intensity` and `P.perimeter`. This implies that λ_3 will be related to any HF that takes either `P.intensity`, `P.perimeter`, or both as inputs. We proceed to build a relational structure among the HFs and DSPs. As shown in Figure 3(b), this structure shows which *independent* DSPs each HF operates over. This relational structure implicitly encodes dependency information about the HFs — if an edge points from one primitive to n HFs, those n HFs are in an n -way relation by virtue of sharing that primitive. This dependency information can more formally be encoded in a factor graph shown in Figure 3(c) and discussed in the next section. Note that we chose a particular programmatic setup for creating DSPs and HFs to explain how static analysis can infer dependencies; however, this process can be modified to work with other setups that encode DSPs and HFs as well.

2.3 Creating the Generative Model

We now describe the generative model used to predict the true class labels. The Coral model uses a factor graph (Figure 3(c)) to model the relationship between the primitives (p),

heuristic functions (λ), and latent class label ($Y \in \{-1, 1\}$). We show that by incorporating information about how primitives are shared across HFs from static analysis, this factor graph infers all dependencies between the heuristics that are guaranteed to be present. In the next section, we describe how Coral recovers additional dependencies between the heuristics by studying empirical relationships between the primitives.

Modeling Heuristic Function Dependencies—With dependencies inferred via static analysis, our goal is to learn the accuracies for each HF and assign labels to training data accordingly. The factor graph thus consists of two types of factors: accuracy factors ϕ^{Acc} and HF factors from static analysis ϕ^{HF} .

The accuracy factors specify the accuracy of each heuristic function λ_i and are defined as

$$\phi_i^{\text{Acc}}(Y, \lambda_i) = Y \lambda_i, \quad i=1, \dots, n$$

where n is the total number of heuristic functions.

The static analysis factors ensure that the heuristics are correctly evaluated based on the HF dependencies found in the previous section. They ensure that a probability of zero is given to any configuration where an HF does not have the correct value given the primitives it depends on. The static analysis factors are defined as

$$\phi_i^{\text{HF}}(\lambda_i, p_1, \dots, p_m) = \begin{cases} 0 & \text{if } \lambda_i \text{ is valid given } p_1, \dots, p_m \\ -\infty & \text{otherwise} \end{cases}, \quad i=1, \dots, n$$

Since these factors are obtained directly from static analysis, *they can be recovered with no data*.

However, we note that static analysis is not sufficient to capture all dependencies required in the factor graph to accurately model the process of generating training labels. Specifically, static analysis can

- i. pick up spurious dependencies among HFs that are not truly dependent on each other, or
- ii. miss key dependencies among HFs that exist due to dependencies among the DSPs in the HFs.

(i) can occur if some λ_A takes as input DSPs p_i, p_j and λ_B takes as input DSPs p_i, p_k , but p_i always has the same value. Although static analysis would pick up that λ_A and λ_B share a primitive and should have a dependency, this may not be true if p_j and p_k are independent.

(ii) can occur if two HFs depend on different primitives, but these primitives happen to always have the same value. In this case, it is impossible for static analysis to infer the dependency between the HFs if the primitives have different names and are generated

independently, as described in Section 2.2. A more realistic example comes from our running example, where we would expect the area and perimeter of the tumor to be related.

To account for both cases, it is necessary to capture the possible dependencies that occur among the DSPs to ensure that the dependencies from static analysis do not misspecify the factor graph. We introduce a factor to account for additional dependencies among the primitives, ϕ^{DSP} . There are many possible choices for this dependency factor, but one simple choice is to model pairwise similarity between the primitives. This allows the dependency factor to be represented as

$$\phi^{\text{DSP}}(p_1, \dots, p_m) = \sum_{i < j} \phi_{ij}^{\text{Sim}}(p_i, p_j), \text{ where } \phi_{ij}^{\text{Sim}}(p_i, p_j) = [p_i = p_j].$$

Finally, with three types of factors, the probability distribution specified by the factor graph is

$$P(y, \lambda_1, \dots, \lambda_n, p_1, \dots, p_m) = \exp \left(\sum_{i=1}^n \theta_i^{\text{Acc}} \phi_i^{\text{Acc}} + \sum_{i=1}^n \phi_i^{\text{HF}} + \sum_{i=1}^m \sum_{j=i+1}^m \theta_{ij}^{\text{Sim}} \phi_{ij}^{\text{Sim}} \right)$$

where θ^{Acc} and θ_{ij}^{Sim} are weights that specify the strength of factors ϕ^{Acc} and ϕ_{ij}^{Sim} .

Inferring Dependencies without Data—The HF factors capture all dependencies among the heuristic functions that are not represented by the ϕ^{DSP} factor. The dependencies represented by the ϕ^{DSP} factor are precisely the dependencies that cannot be inferred via static analysis due to the fact that this factor depends solely on the content of the primitives. It is therefore impossible to determine what this factor is without data.

While assuming that we have the true ϕ^{DSP} seems like strong condition, we find that in real-world experiments, including the ϕ^{DSP} factor rarely leads to improvements over the case when we only include the ϕ^{Acc} and ϕ^{HF} factors. In some of our experiments (see Section 3), we use bounding box location, size and object labels as domain-specific primitives for image and video querying tasks. Since these primitives are not correlated, modeling the primitive dependency does not lead to any improvement over just modeling HF dependencies from static analysis. Moreover, in other experiments in which modeling the relation among primitives helps, we observe relatively small benefits (up to 1.1 points) above what modeling HF dependencies provides (up to 3.47 points). Therefore, without any data, it is possible to model the most important dependencies among HFs that lead to significant gains over the case where no dependencies are modeled.

2.4 Generating Probabilistic Training Labels

Given the probability distribution of the factor graph, our goal is to learn the proper weights θ_i^{Acc} and θ_{ij}^{Sim} . Coral adopts structure learning approaches described in recent work [35, 2] that learns dependency structures in the weak supervision setting and maximizes the ℓ_1 -

regularized marginal pseudolikelihood of each primitive to learn the weights of the relevant factors.

To learn the weights of the generative model, we use contrastive divergence [20] as a maximum likelihood estimation routine and maximize the marginal likelihood of the observed primitives. Gibbs sampling is used to estimate the intractable gradients, which are then used in stochastic gradient descent. Because the HFs are typically deterministic functions of the primitives (represented as the $-\infty$ value of the correctness factors for invalid heuristic values), standard Gibbs sampling will not be able to mix properly. As a result, we modify the Gibbs sampler to simultaneously sample one primitive along with all heuristics that depend on it. Despite the fact that the true class label is latent, this process still converges to the correct parameter values [34]. Additionally, the amount of data necessary to learn the parameters scales quasilinearly with the number of parameters. In our case, the number of parameters is just the number of heuristics n and number of relevant primitive similarity dependencies s . We now formally state the conditions for this result, which match those of Ratner et al. [34], and provide the sample complexity of our method. First, we assume that there exists some feasible parameter set $\Theta \subset \mathbb{R}^n$ that is known to contain the parameter $\theta^* = (\theta^{\text{Acc}}, \theta^{\text{Sim}})$ that models the true distribution π^* of the data:

$$\exists \theta^* \in \Theta \text{ s. t. } \forall \pi^*(p_1, \dots, p_m, Y) = \mu_{\theta}(p_1, \dots, p_m, Y). \quad (1)$$

Next, we must be able to accurately learn θ^* if we are provided with labeled samples of the true distribution. Specifically, there must be an asymptotically unbiased estimator $\hat{\theta}$ that takes some set of labeled data T independently sampled from π^* such that for some $c > 0$,

$$\text{Cov}(\hat{\theta}(T)) \leq (2c|T|)^{-1} I. \quad (2)$$

Finally, we must have enough sufficiently accurate heuristics so that we have a reasonable estimate of Y . For any two feasible models $\theta_1, \theta_2 \in \Theta$,

$$\mathbb{E}_{(p_1, \dots, p_m, Y) \sim \mu_{\theta_1}} \left[\text{Var}_{(p'_1, \dots, p'_m, Y') \sim \mu_{\theta_2}} (Y' | p_1 = p'_1, \dots, p_m = p'_m) \right] \leq \frac{c}{n+s} \quad (3)$$

Proposition 1: *Suppose that we run stochastic gradient descent to produce estimates of the weights $\hat{\theta} = (\hat{\theta}^{\text{Acc}}, \hat{\theta}^{\text{Sim}})$ in a setup satisfying conditions (1), (2), and (3). Then, for any fixed error $\epsilon > 0$, if the number of unlabeled data points N is at least $\Omega[(n+s) \log(n+s)]$, then our expected parameter error is bounded by $\mathbb{E}[\|\hat{\theta} - \theta^*\|^2] \leq \epsilon^2$.*

The proof follows from the sample complexity of Ratner et al. [34] and appears in the appendix. With the weights $\hat{\theta}_i^{\text{Acc}}$ and $\hat{\theta}_{ij}^{\text{Sim}}$ maximizing the marginal likelihood of the observed primitives, we have a fully specified factor graph and complete generative model,

which can be used to predict the latent class label. For each data point, we have the domain-specific primitives, from which the heuristic functions can compute noisy labels. Through the accuracy factors, we then estimate a distribution for the latent class label and use these noisy labels to train a discriminative model.

In Figure 4, we present a simulation to empirically compare our sample complexity with that of structure learning [2]. In our simulation, we have n HFs, each with an accuracy of 75%, and explore settings with where there exists one binary, 3-ary and 4-ary dependency among the HFs. The dependent HFs share exactly one primitive, and the primitives themselves are independent ($s = 0$). As N increases, both methods improve in performance due to improved estimates of the heuristic accuracies and dependencies. In the case with a binary dependency, structure learning recovers the necessary dependency with few samples, and has similar performance to Coral. In contrast, in the second and third settings with high-order dependencies, structure learning struggles to recover the relevant dependency, and consistently performs worse than Coral even as more training data is provided.

3 Experimental Results

We seek to experimentally validate the following claims about our approach. Our first claim is that HF dependencies inferred via static analysis (HF Dep) improve significantly over a model that does not take dependencies into account. Second, we compare to a structure learning approach for weak supervision [2] and show how we outperform it over a variety of domains. Finally, we show that in case there are primitive dependencies, Coral can learn and model those as well (HF+DSP Dep). Our classification tasks range from specialized medical domains to natural images and video, and we include details of the DSPs and HFs in the Appendix. We compare our approach to generative models that learn the accuracies of different heuristics, specifically one that assumes the heuristics are independent (Indep), and Bach et al. [2] that *learns* the inter-heuristic dependencies (Learn Dep). We also compare to majority vote (MV) and the fully supervised (FS) case, and measure the performance of the discriminative model trained with labels generated using the above methods.

Visual Genome and ActivityNet Classification—We explore how to extract complex relations in images and videos given object labels and their bounding boxes. We used subsets of two datasets, Visual Genome [26] and ActivityNet [11], and defined our task as finding images of “a person biking down a road” and finding basketball videos, respectively. For both tasks, a small set of DSPs were shared heavily among HFs, and modeling the dependencies observed by static analysis led to a significant improvement over the independent case. Since these dependencies involved groups of 3 or more heuristics, Coral improved significantly, by up to 3.81 F1 points, over structure learning as well, which was unable to model these dependencies due to the lack of enough data. Moreover, modeling primitive dependencies did not help since the primitives were indeed independent (Table 1). We report our results for these tasks in terms of the F1 score, harmonic mean of the precision and recall.

Bone Tumor Classification—We used a set of 802 labeled bone tumor X-ray images along with their radiologist-drawn segmentations. We defined our task to differentiate

between aggressive and non-aggressive tumors, generated HFs that were a combination of hand-tuned rules and decision-tree generated rules (tuned on a small subset of the dataset). The discriminative model utilized a set of 400 hand-tuned features (note that there is no overlap between these features and the DSPs) that encoded various shape, texture, edge and intensity-based characteristics. Although there were no explicitly shared primitives in this dataset, the generative model was still able to model the training labels more accurately with knowledge of *how many primitives* the heuristics operated over, thus improving over the independent case significantly. Moreover, a small dataset size hindered structure learning, which gave a minimal boost over the independent case (Table 1). When we used heuristics in Coral to label an additional 800 images that had no ground truth labels, we beat the previous FS score by 3.07 points (Figure 5, Table 1).

Mammogram Tumor Classification—We used the DDSM-CBIS [39] dataset, which consists of 1800 scanned film mammograms and associated segmentations for the tumors in the form of binary masks. We defined our task to identify a tumor as malignant or benign. Each heuristic only operated over one primitive, resulting in no dependencies that static analysis could identify. In this case, structure learning performed better than Coral when we only used static analysis to infer dependencies (Figure 5). However, including primitive dependencies allowed us to match structure learning, resulting in a 1.11 point improvement over the independent case (Table 1).

4 Related Work

As the need for labeled training data grows, a common alternative is to utilize weak supervision sources such as distant supervision [12, 30], multi-instance learning [37, 21], and heuristics [10, 42]. Specifically for images, weak supervision using object detection and segmentation or visual databases is a popular technique as well (detailed discussion in Appendix). Estimating the accuracies of these sources without access to ground truth labels is a classic problem [15]. Methods such as crowdsourcing, boosting, co-training, and learning from noisy labels are some of the popular approaches that can combine various sources of weak supervision to assign noisy labels to data (detailed discussion in Appendix). However, Coral does not require *any* labeled data to model the dependencies among the domain-specific primitives and heuristics, which can be interpreted as workers, classifiers or views for the above methods.

Recently, generative models have also been used to combine various sources of weak supervision [1, 44, 38]. One specific example, data programming [34], proposes using multiple sources of weak supervision for text data in order to describe a generative model and subsequently learns the accuracies of these source. Coral also focuses on multiple programmatically encoded heuristics that can weakly label data and learns their accuracies to assign labels to training data. However, Coral adds an additional layer of domain-specific primitives in its generative model, which allows it to generalize beyond text-based heuristics. It also *infers* the dependencies among the heuristics and the primitives, rather than requiring users to specify them. Other previous work also assume that this structure in generative models is user-specified [1, 44, 38]. Recently, Bach et al. [2] showed that it is possible to learn the dependency structure among sources of weak supervision with a sample

complexity that scales sublinearly with the number of possible pairwise dependencies. Coral instead identifies the dependencies among the heuristic functions by inspecting the content of the programmable functions, therefore relying on significantly less data to learn the generative model structure. Moreover, Coral can also pick up higher-order dependencies, for which structured learning needs large amounts of data to detect.

5 Conclusion and Future Work

In this paper, we introduced Coral, a paradigm that models the dependency structure of weak supervision heuristics and systematically combines their outputs to assign probabilistic labels to training data. We described how Coral takes advantage of the programmatic nature of these heuristics in order to infer dependencies among them via static analysis and requires a sample complexity that is quasilinear in the number of heuristics and relations found. We showed how Coral leads to significant improvements in discriminative model accuracy over traditional structure learning approaches across various domains. Coral scratches the surface of the possible ways weak supervision can borrow from the field of programming languages, especially as they are applied to large magnitudes of data and need to be encoded programmatically. We look at a natural extension of treating the process of encoding heuristics as writing functions and hope to further explore the interactions between systematic training set creation and concepts from the programming language field.

Acknowledgments

We thank Shoumik Palkar, Stephen Bach, and Sen Wu for their helpful conversations and feedback. We are grateful to Darvin Yi for his assistance with the DDSM dataset based experiments and associated deep learning models. We acknowledge the use of the bone tumor dataset annotated by Drs. Christopher Beaulieu and Bao Do and carefully collected over his career by the late Henry H. Jones, M.D. (aka “Bones Jones”).

This material is based on research sponsored by Defense Advanced Research Projects Agency (DARPA) under agreement number FA8750-17-2-0095. We gratefully acknowledge the support of the DARPA SIMPLEX program under No. N66001-15-C-4043, DARPA FA8750-12-2-0335 and FA8750-13-2-0039, DOE 108845, the National Science Foundation (NSF) Graduate Research Fellowship under No. DGE-114747, Joseph W. and Hon Mai Goodman Stanford Graduate Fellowship, the Moore Foundation, National Institute of Health (NIH) U54EB020405, the Office of Naval Research (ONR) under awards No. N000141210041 and No. N000141310129, the Moore Foundation, the Okawa Research Grant, American Family Insurance, Accenture, Toshiba, and Intel. This research was supported in part by affiliate members and other supporters of the Stanford DAWN project: Intel, Microsoft, Teradata, and VMware.

The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA, AFRL, NSF, NIH, ONR, or the U.S. government.

References

1. Alfonseca, E., Filippova, K., Delort, J-Y., Garrido, G. Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2. Association for Computational Linguistics; 2012. Pattern learning for relation extraction with a hierarchical topic model; p. 54-59.
2. Bach, SH., He, B., Ratner, A., Ré, C. Learning the structure of generative models without labeled data. ICML; 2017.

3. Balsubramani A, Freund Y. Scalable semi-supervised aggregation of classifiers. *Advances in Neural Information Processing Systems*. 2015:1351–1359.
4. Banerjee I, Hahn L, Sonn G, Fan R, Rubin DL. Computerized multiparametric mr image analysis for prostate cancer aggressiveness-assessment. 2016 arXiv preprint arXiv:1612.00408.
5. Barb AS, Shyu C-R, Sethi YP. Knowledge representation and sharing using visual semantic modeling for diagnostic medical image databases. *IEEE Transactions on Information Technology in Biomedicine*. 2005; 9(4):538–553. [PubMed: 16379371]
6. Blaschko M, Vedaldi A, Zisserman A. Simultaneous object detection and ranking with weak supervision. *Advances in neural information processing systems*. 2010:235–243.
7. Blum, A., Mitchell, T. Combining labeled and unlabeled data with co-training. *Proceedings of the eleventh annual conference on Computational learning theory*; ACM; 1998. p. 92-100.
8. Branson, S., Perona, P., Belongie, S. Strong supervision from weak annotation: Interactive training of deformable part models. *Computer Vision (ICCV), 2011 IEEE International Conference on*; IEEE; 2011. p. 1832-1839.
9. Brooks, RA., Creiner, R., Binford, TO. The acronym model-based vision system. *Proceedings of the 6th international joint conference on Artificial intelligence-Volume 1*; Morgan Kaufmann Publishers Inc; 1979. p. 105-113.
10. Bunescu, R., Mooney, R. Learning to extract relations from the web using minimal supervision. *ACL*; 2007.
11. Caba Heilbron, F., Escorcia, V., Ghanem, B., Carlos Niebles, J. Activitynet: A large-scale video benchmark for human activity understanding. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; 2015. p. 961-970.
12. Craven, M., Kumlien, J., et al. Constructing biological knowledge bases by extracting information from text sources. *ISMB*; 1999. p. 77-86.
13. Dai, J., He, K., Sun, J. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. *Proceedings of the IEEE International Conference on Computer Vision*; 2015. p. 1635-1643.
14. Dalvi, N., Dasgupta, A., Kumar, R., Rastogi, V. Aggregating crowdsourced binary ratings. *Proceedings of the 22nd international conference on World Wide Web*; ACM; 2013. p. 285-294.
15. Dawid AP, Skene AM. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Applied statistics*. 1979:20–28.
16. Fischler MA, Elschlager RA. The representation and matching of pictorial structures. *IEEE Transactions on computers*. 1973; 100(1):67–92.
17. Fukui A, Park DH, Yang D, Rohrbach A, Darrell T, Rohrbach M. Multimodal compact bilinear pooling for visual question answering and visual grounding. 2016 arXiv preprint arXiv:1606.01847.
18. Haralick RM, Shanmugam K, et al. Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics*. 1973; 3(6):610–621.
19. Hearst, MA. Automatic acquisition of hyponyms from large text corpora. *Proceedings of the 14th conference on Computational linguistics-Volume 2*; Association for Computational Linguistics; 1992. p. 539-545.
20. Hinton GE. Training products of experts by minimizing contrastive divergence. *Neural computation*. 2002; 14(8):1771–1800. [PubMed: 12180402]
21. Hoffmann, R., Zhang, C., Ling, X., Zettlemoyer, L., Weld, DS. Knowledge-based weak supervision for information extraction of overlapping relations. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*; Association for Computational Linguistics; 2011. p. 541-550.
22. Joglekar, M., Garcia-Molina, H., Parameswaran, A. Comprehensive and reliable crowd assessment algorithms. *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*; IEEE; 2015. p. 195-206.
23. Kang, D., Emmons, J., Abuzaid, F., Bailis, P., Zaharia, M. Optimizing deep cnn-based queries over video streams at scale. *CoRR*. abs/1703.02529, 2017. URL <http://arxiv.org/abs/1703.02529>

24. Karpathy, A., Fei-Fei, L. Deep visual-semantic alignments for generating image descriptions. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2015. p. 3128-3137.
25. Kaus MR, Warfield SK, Nabavi A, Black PM, Jolesz FA, Kikinis R. Automated segmentation of mr images of brain tumors 1. Radiology. 2001; 218(2):586–591. [PubMed: 11161183]
26. Krishna R, Zhu Y, Groth O, Johnson J, Hata K, Kravitz J, Chen S, Kalantidis Y, Li L-J, Shamma DA, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. 2016 arXiv preprint arXiv:1602.07332.
27. Kurtz C, Depeursinge A, Napel S, Beaulieu CF, Rubin DL. On combining image-based and ontological semantic dissimilarities for medical image retrieval applications. Medical image analysis. 2014; 18(7):1082–1100. [PubMed: 25036769]
28. Lu, C., Krishna, R., Bernstein, M., Fei-Fei, L. Visual relationship detection with language priors. European Conference on Computer Vision; Springer; 2016. p. 852-869.
29. Meinshausen N, Bühlmann P. High-dimensional graphs and variable selection with the lasso. The annals of statistics. 2006:1436–1462.
30. Mintz, M., Bills, S., Snow, R., Jurafsky, D. Distant supervision for relation extraction without labeled data. 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; Association for Computational Linguistics; 2009. p. 1003-1011.
31. Oliver A, Freixenet J, Marti J, Pérez E, Pont J, Denton ER, Zwiggelaar R. A review of automatic mass detection and segmentation in mammographic images. Medical image analysis. 2010; 14(2): 87–110. [PubMed: 20071209]
32. Oquab, M., Bottou, L., Laptev, I., Sivic, J. Is object localization for free? - Weakly-supervised learning with convolutional neural networks. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2015. p. 685-694.
33. Raicu DS, Varutbangkul E, Furst JD, Armato SG III. Modelling semantics from image data: opportunities from lidc. International Journal of Biomedical Engineering and Technology. 2009; 3(1–2):83–113.
34. Ratner AJ, De Sa CM, Wu S, Selsam D, Ré C. Data programming: Creating large training sets, quickly. Advances in Neural Information Processing Systems. 2016:3567–3575.
35. Ravikumar P, Wainwright MJ, Lafferty JD, et al. High-dimensional ising model selection using l1-regularized logistic regression. The Annals of Statistics. 2010; 38(3):1287–1319.
36. Redmon, J., Divvala, S., Girshick, R., Farhadi, A. You only look once: Unified, real-time object detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2016. p. 779-788.
37. Riedel, S., Yao, L., McCallum, A. Modeling relations and their mentions without labeled text. Joint European Conference on Machine Learning and Knowledge Discovery in Databases; Springer; 2010. p. 148-163.
38. Roth, B., Klakow, D. Combining generative and discriminative model scores for distant supervision. EMNLP; 2013. p. 24-29.
39. Sawyer-Lee R, Gimenez F, Hoogi A, Rubin D. Curated breast imaging subset of ddsd. 2016
40. Schapire, RE., Freund, Y. Boosting: Foundations and algorithms. MIT press; 2012.
41. Sharma N, Aggarwal LM, et al. Automated medical image segmentation techniques. Journal of medical physics. 2010; 35(1):3. [PubMed: 20177565]
42. Shin J, Wu S, Wang F, De Sa C, Zhang C, Ré C. Incremental knowledge base construction using DeepDive. Proceedings of the VLDB Endowment. 2015; 8(11):1310–1321. [PubMed: 27144081]
43. Stewart, R., Ermon, S. Label-free supervision of neural networks with physics and domain knowledge. AAAI; 2017.
44. Takamatsu, S., Sato, I., Nakagawa, H. Reducing wrong labels in distant supervision for relation extraction. Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1; Association for Computational Linguistics; 2012. p. 721-729.
45. Xia, W., Domokos, C., Dong, J., Cheong, L-F., Yan, S. Semantic segmentation without annotating segments. Proceedings of the IEEE International Conference on Computer Vision; 2013. p. 2176-2183.

46. Yi D, Zhou M, Chen Z, Gevaert O. 3-d convolutional neural networks for glioblastoma segmentation. 2016 arXiv preprint arXiv:1611.04534.
47. Zhang Y, Chen X, Zhou D, Jordan MI. Spectral methods meet em: A provably optimal algorithm for crowdsourcing. *Journal of Machine Learning Research*. 2016; 17(102):1–44.
48. Zhao P, Yu B. On model selection consistency of lasso. *Journal of Machine learning research*. Nov. 2006 7:2541–2563.

Appendix

A Extended Related Works

We provide additional details here of how writing heuristics over primitives is a popular technique, making Coral widely applicable to existing methods as well. We then discuss methods that are related to Coral in combining various sources of weak supervision to generate noisy labels for the data at hand.

Weak Supervision over Primitives

A popular approach for creating training sets is to provide weak or distant supervision to label data based on information from a knowledge base [12, 30], crowdsourcing [15], heuristic patterns [10, 19], user input [42], a set of user-defined labeling functions [34], or hand-engineered constraints [43]. Our inspiration for Coral came from observing various weak supervision and image description techniques developed in the field of computer vision. Early work looked at describing images in terms of a set of primitives to find instances of described objects in images [16, 9]. More recently, learned characteristics such as bounding boxes from object detection and areas from image segmentation have been used in order to weakly supervise more complex image-based learning tasks [13, 45, 6, 32, 8]. Moreover, the recent development of a ‘knowledge base’ for images, Visual Genome [26], has provided access to a image database with rich, crowdsourced attribute and relational information. This data in turn powers other methods that rely on the information in Visual Genome to supervise other tasks [28, 17]. This trend of using the possibly noisy information about the data at hand in order to weakly supervise models that look for complex relationships is one of the main motivations for Coral, which is a domain-agnostic method that can combine various sources of weak supervision while modeling the relationship among the sources as well as the primitives they operate over.

A similar methodology of using interpretable characteristics of images for classification tasks exists in the medical field as well [5, 33]. Other techniques such as content-based image retrieval (CBIR) rely on low-level, quantitative features of images, such as tumor texture, obtained from image analysis technique to query similar images in a database of radiology images. However, there remains a gap between these content-based primitives and semantic, interpretable primitives that humans use in order to describe these images [27]. Coral could combine the two sources of supervision in order to build a system that uses the two methodologies in an optimal manner, while also helping discover correlations among the semantic and quantitative primitives.

Generative Models to Generate Training Labels

Crowdsourcing is another area where the estimating worker accuracy without hand-labeled data is a well-studied problem [14, 22, 47]. Boosting is also a related method that combines the output of multiple weak classifiers to create a stronger classifier [40], and recent work has further looked at leveraging unlabeled data [3]. Co-training is another process that uses a small labeled dataset and a large amount of unlabeled data by selecting two conditionally independent views of the data [7].

B Additional Experimental Results

We present the results from Section 3 in a table format here. We also provide a list of the domain-specific primitives and heuristic functions used in each of our experiments, with a short description of what they mean.

Table 2

Numbers for Figure 5(*ed have F1 scores, rest are accuracy). MV: majority vote; Indep: assume heuristics independent; Learn: model learned dependencies; HF-Dep: model heuristic dependencies from static analysis; HF+DSP Dep: also model primitive dependencies; FS: fully supervised

Application	MV	Indep	Learn Dep	HF Dep	HF+DSP Dep	FS
Visual Genome*	42.02	46.61	46.61	49.51	49.51	50.25
ActivityNet*	46.44	48.86	48.86	52.67	52.67	54.54
Bone Tumor	65.72	67.32	67.83	70.79	70.89	67.82
Mammogram	63.8	67.31	68.42	67.31	68.42	69.06

C Proof of Proposition 1

Proposition 1

Suppose that we run stochastic gradient descent to produce estimates of the weights $\hat{\theta} = (\hat{\theta}^{Acc}, \hat{\theta}^{Sim})$ in a setup satisfying conditions (1), (2), and (3). Then, for any fixed error $\epsilon > 0$, if the number of unlabeled data points N is at least $\Omega[(n + s) \log(n + s)]$, then our expected parameter error is bounded by $\mathbb{E}[\|\hat{\theta} - \theta^\|^2] \leq \epsilon^2$.*

Proof—The proof of proposition 1 closely follows Theorem 2 of Ratner et al. [34]. First, notice that all of the necessary conditions of this theorem are satisfied. The number of weights to be learned in this theorem is M . In our setting, $M = n + s$. Notice that if we have at least $\Omega[(n + s) \log(n + s)]$ unlabeled data points, then we satisfy the conditions of the theorem. As a result, the bound on the expected parameter error directly follows.

Table 3

Domain-specific Primitives for Mammogram Tumor Classification

Primitives	Type	Description
area	float	area of the tumor
diameter	float	diameter of the tumor
eccentricity	float	eccentricity of the image
perimeter	float	perimeter of the tumor
max_intensity	float	maximum intensity of the tumor
mean_intensity	float	mean intensity of the tumor

Table 4

Domain-specific Primitives for Image Querying

Primitives	Type	Description
person	bool	image contains a person/man/woman
person.position	(float, float)	coordinates of the person
person.area	float	area of bounding box of person
road	bool	image contains a road/street
car	bool	image contains car/bus/truck
bike	bool	image contains a bike/bicycle/cycle
bike.position	(float, float)	coordinates of the bike
bike.area	float	area of bounding box of bike

Table 5

Domain-specific Primitives for Video Classification

Primitives	Type	Description
person	bool	image contains a person
ball	bool	image contains sports ball
(person.top, person.bottom)	(float, float)	top and bottom coordinates of the person
(ball.top, ball.bottom)	(float, float)	top and bottom coordinates of the ball
ball.color	(float, float, float)	R,G,B colors of the ball
vertical_distance	float	cumulative absolute difference in ball.top values over frames

Table 6
Domain-specific Primitives for Bone Tumor Classification

Primitives	Type	Description
daube_hist_164	float	Daubechies features
daube_hist_224	float	Daubechies features
daube_hist_201	float	Daubechies features
window_std	float	
window_median	float	
scale_median	float	
lesion_density	float	quantify edge sharpness along the lesion contour
edge_sharpness	float	quantify edge sharpness along the lesion contour
equiv_diameter	float	describe the morphology of the lesion
area	float	describe the morphology of the lesion
perimeter	float	describe the morphology of the lesion
area_perimeter_ratio	float	ratio of area and perimeter
shape_solidity	float	
laplacian_entropy	float	Laplacian energy features
sobel_entropy	float	Sobel energy features
glcm_contrast	float	capture occurrence of gray level pattern within the lesion
glcm_homogeneity	float	capture occurrence of gray level pattern within the lesion
histogram_egde	float	quantify edge sharpness along the lesion contour
mean_diff_in_out	float	

Table 7
Heuristic Functions for Mammogram Tumor Classification

Name	Heuristic Function	Description
hf_area	1 if area \geq 100000; -1 if area \leq 30000	Large tumor area indicates malignant tumors
hf_diameter	1 if diameter \geq 400; -1 if diameter \leq 200	High diameter indicates malignant tumors
hf_eccentricity	1 if eccentricity \leq 0.4; -1 if eccentricity \geq 0.6	Low eccentricity indicates malignant tumors
hf_perimeter	1 if perimeter \geq 4000; -1 if perimeter \leq 2500	High perimeter indicates malignant tumors
hf_max_intensity	1 if max_intensity \geq 70000; -1 if max_intensity \leq 50000	High maximum intensity indicates malignant tumors
hf_mean_intensity	1 if mean_intensity \geq 45000; -1 if mean_intensity \leq 30000	High mean intensity indicates malignant tumors

Table 8

Heuristic Functions for Image Querying

Name	Heuristic Function	Description
hf_street	1 if person and road; -1 if person and ! road; 0 otherwise	Indicates if a street is present when person is present, doesn't assign a label when person is not present
hf_vehicles	1 if person and car; -1 if person and ! car; 0 otherwise	Indicates if a car is present when person is present, doesn't assign a label when person is not present
hf_positions	-1 if !bike or !person; 1 if (person.position-bike.position)<=1; 0 otherwise	Indicates if person and bike are close in the image
hf_size	0 if !bike or !person; -1 if (person.area-bike.area)>=1000; 1 otherwise	Indicates if the difference in area of the bike and person is less than a threshold
hf_number	-1 if !bike or !person; 1 if num_persons=num_bikes; -1 if num_persons<num_bikes; 0 otherwise	Indicates if number of persons and bikes are equal

Table 9

Heuristic Functions for Video Classification

Name	Heuristic Function	Description
hf_person_ball	1 if person and ball; -1 otherwise	Indicates if person and sports ball were present in any frame of the video
hf_distance	-1 if person.top-ball.bottom>=2; 1 if person.top-ball.bottom<=1; 0 otherwise	Indicates if the distance between person and sports ball is less than a threshold
hf_ball_color	1 if ball and ball.color-basketball.color <= 80; -1 otherwise	Indicates if the color of the ball is similar to the color of a basketball
hf_temporal	1 if vertical_distance >=15; -1 otherwise	Indicates if sufficient vertical distance was covered by the ball over frames

Table 10

Heuristic Functions for Bone Tumor Classification

Name	Heuristic Function
hf_daube	1 if histogram_164<0.195545 and histogram_224<-0.469812; -1 if histogram_164<0.195545 and histogram_201<0.396779; 1 if histogram_164 <0.195545; -1 otherwise
hf_edge	-1 if window_std<-0.0402606; -1 if window_median<-0.544591; -1 if scale_median<-0.512551; 1 otherwise
hf_lesion	-1 if lesion_density<0; 1 if lesion_density>1 and edge_sharpness<0; 1 if lesion_density>5 and edge_sharpness>-1; 0 otherwise
hf_shape	-1 if equiv_diameter<-0.3; -1 if equiv_diameter>0 and area_perimeter_ratio<0.5 and shape_solidity<0.1; 1 if equiv_diameter >0 and area_perimeter_ratio<0.5 and

Name	Heuristic Function
	shape_solidity>0.75; 1 if equiv_diameter>0 and area_perimeter_ratio>1; 0 otherwise
hf_sobel_laplacian	-1 if laplacian_entropy<0.2; 1 if laplacian_entropy>0.4 and sobel_entropy<-0.75; 1 if laplacian_entropy>0.4 and sobel_entropy >-0; 0 otherwise
hf_glcm	-1 if glcm_contrast<0.15 and glcm_homogeneity<0; 1 if glcm_contrast<0.15 and glcm_homogeneity>0.5; -1 if glcm_contrast>0.25; 0 otherwise
hf_first_order	-1 if histogram_egde<0.5; 1 if histogram_egde>-0.3 and mean_diff_in_out <-0.75; 1 if histogram_egde>-0.3 and mean_diff_in_out>-0.5; 0 otherwise

We omit the descriptions of heuristic functions for bone tumor classification due to their complexity.




Figure 1. Running example for Coral paradigm. Users apply standard algorithms to segment tumors from the X-ray and extract the *domain-specific primitives* from the image and segmentation. They write *heuristic functions* over the primitives that output a noisy label for each image. The generative model takes these as inputs and provides probabilistic training labels for the discriminative model.




Figure 2. (a) The `create_primitives` function that generates primitives. (b) Part of the AST for the `create_primitives` function. (c) The composition structure that results from traversing the AST.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

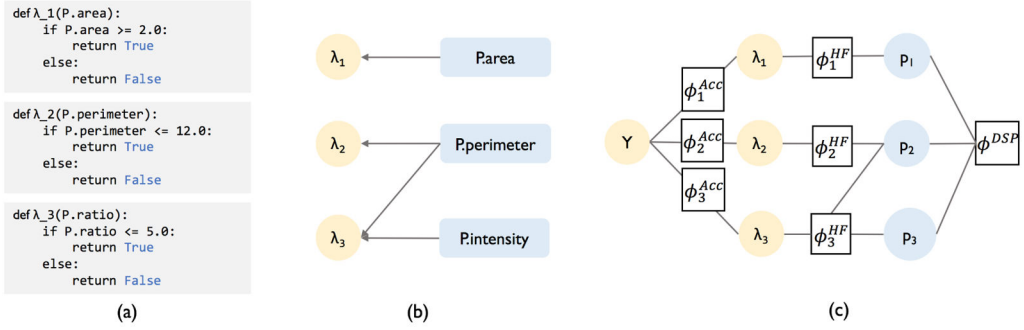


Figure 3.

(a) shows the encoded HFs. (b) shows the HF dependency structure where DSP nodes have an edge going to the HFs that use them as inputs (explicitly or implicitly). (c) shows the factor graph Coral uses to model the relationship between HFs, DSPs, and latent class label Y .

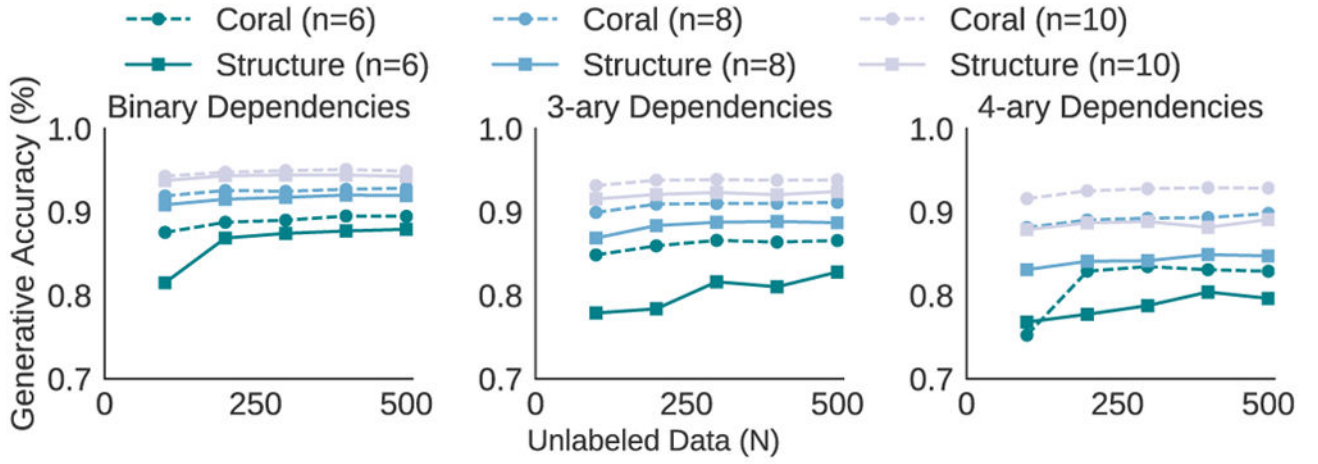


Figure 4. Simulation comparing the generative model accuracy with structure learning and Coral.




Figure 5. Discriminative model performance comparing HF Dep (HF dependencies from static analysis) and HF+DSP Dep (HF and DSP dependencies) to other methods. Numbers in Appendix.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Heuristic Function (HF) and Domain-Specific Primitive (DSP) statistics. Discriminative model improvement with HF+DSP Dep over other methods.

Table 1

Application	Statistics			Improvement Over			
	# DSPs	# HFs	# Shared DSPs	MV	Indep	Learn Dep	FS
Visual Genome	7	5	2	7.49 *	2.90 *	2.90 *	-0.74 *
ActivityNet	5	4	2	6.23 *	3.81 *	3.81 *	-1.87 *
Bone Tumor	17	7	0	5.17	3.57	3.06	3.07
Mammogram	6	6	0	4.62	1.11	0	-0.64

* improvements shown in terms of F1 score, rest in terms of accuracy. ActivityNet model is LR using VGGNet embeddings as features.