
Co-training Improves Prompt-based Learning for Large Language Models

Hunter Lang¹ Monica Agrawal¹ Yoon Kim¹ David Sontag¹

Abstract

We demonstrate that co-training (Blum & Mitchell, 1998) can improve the performance of prompt-based learning by using *unlabeled data*. While prompting has emerged as a promising paradigm for few-shot and zero-shot learning, it is often brittle and requires much larger models compared to the standard supervised setup. We find that co-training makes it possible to improve the original prompt model and at the same time learn a smaller, downstream task-specific model. In the case where we only have partial access to a prompt model (e.g., output probabilities from GPT-3 (Brown et al., 2020)) we learn a calibration model over the prompt outputs. When we have full access to the prompt model’s gradients but full finetuning remains prohibitively expensive (e.g., T0 (Sanh et al., 2022)), we learn a set of soft prompt continuous vectors to iteratively update the prompt model. We find that models trained in this manner can significantly improve performance on challenging datasets where there is currently a large gap between prompt-based learning and fully-supervised models.

1. Introduction

Prompt-based learning, in which a pretrained language model is adapted to various tasks by priming on natural language prompts, has emerged as a promising framework for few-shot and zero-shot learning (Brown et al., 2020; Liu et al., 2021a; Wei et al., 2021; Sanh et al., 2022). While intriguing, these methods can be sensitive to trivial cosmetic artifacts, including variations in prompt wording and the ordering of examples (Lu et al., 2021; Zhao et al., 2021; Kumar & Talukdar, 2021). Further, the models used in prompt-based learning (e.g., GPT-3, T0) are much larger than those typically used for standard fine-tuning. These factors make prompt-based learning difficult to use in practice.

Given a small amount of labeled data, one could evaluate

¹MIT CSAIL. Correspondence to: <hjl@mit.edu>.

the performance of each prompt and re-calibrate the prompt outputs to improve performance. However, (i) this reliance on labeled data goes against the goal of few-shot learning, and (ii) even with oracle calibration, some prompts have sub-par accuracy. Recently, to address issue (i), Zhao et al. (2021) developed a *data-free* calibration method that can dramatically improve the accuracy of few-shot prompts for GPT-3. We build on their work by showing how to use *unlabeled data* to further improve performance.

To leverage unlabeled data, we use *co-training* (Blum & Mitchell, 1998), which operates on two views of each data point X : $\phi_0(X)$ and $\phi_1(X)$. For example, in a clinical diagnosis system, $\phi_0(X)$ could be laboratory test results and $\phi_1(X)$ an X-ray image. A pair of models (h_0 and h_1 respectively) takes turns labeling a large unlabeled training set, and each model is trained on the confident pseudo-labels from the other. Model h_0 only uses $\phi_0(X)$, and model h_1 uses $\phi_1(X)$. By using complementary information in the views ϕ_0, ϕ_1 and the different inductive biases from models h_0, h_1 , co-training allows each model to learn from the other without labeled data. The initial signal to start the co-training process is provided by a “guess” at a model h_0 . To combine co-training and prompt-based learning, we use outputs from a large prompt-based model as $\phi_0(X)$ and the pre-trained representation from a much smaller language model (e.g., DeBERTa (He et al., 2021)) as $\phi_1(X)$. We specify the models h_0 and h_1 based on whether we have partial access to the prompt model (querying GPT-3) or full access (locally training T0).

In *partial access*, we only have access to the large model’s output probabilities. In this case, we use unlabeled data to learn a model h_0 that both calibrates individual prompts and ensembles multiple prompts. We refer to this as the *label model*. We use Calibrate-Before-Use (Zhao et al., 2021) to initialize the calibration parameters of this model for each prompt, and we initialize the ensembling parameters to approximate majority vote. We then refine this initial guess for h_0 with co-training. We use the pre-trained representation from DeBERTa (He et al., 2021) for $\phi_1(X)$ and train the last few layers of that model as h_1 . The only labeled data used is the set of k examples used in the input prompts. Figure 1 (left) shows the co-training process for partial access.

We also study a *full access* setting using T0 (Sanh et al.,

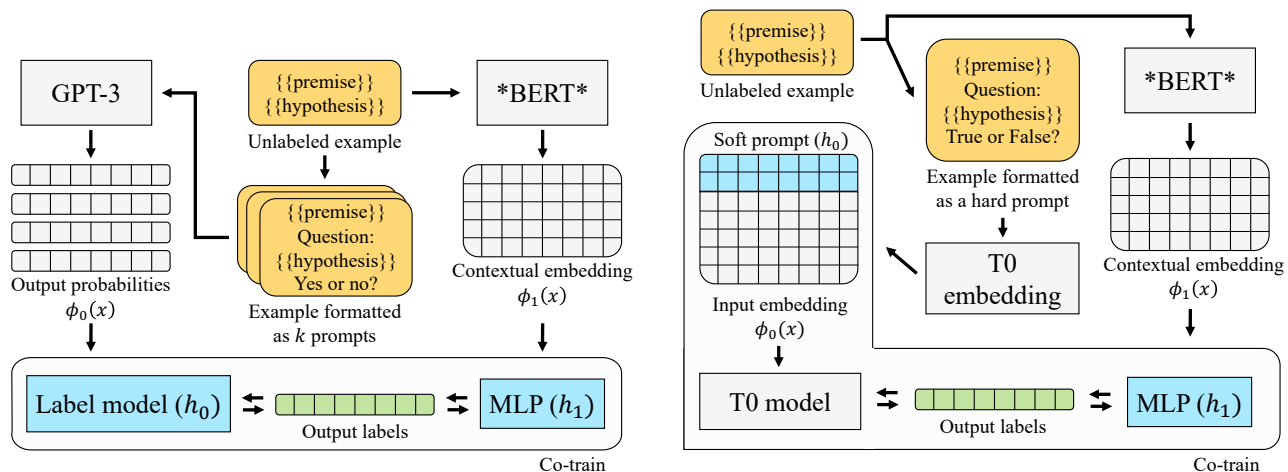


Figure 1. The setup for our two applications of co-training to prompting for a binary entailment classification dataset (RTE). Parameters in blue are trainable; models in gray are fixed. Left: training a “label model” for post-hoc calibration and ensembling of multiple prompts. Here the prompts and the model (GPT-3) are fixed, and we co-train the calibration / ensembling parameters with the task-specific model (e.g., DeBERTa). Right: training a soft prompt. Here the input is encoded as a hard prompt and the embedding matrix of the input sequence is obtained. A $L \times d$ matrix of trainable parameters (the “soft prompt”) is prepended to this embedding, and the combined embedding sequence is passed through T0 to get output predictions. We co-train the soft prompt with the view 1 model (e.g., DeBERTa).

2022) instead of GPT-3, so we can introspect the large prompt model. We derive the view $\phi_0(X)$ and the model h_0 from T0¹. However, instead of fully fine-tuning T0 during co-training, we focus on *soft prompt tuning*, which trains several orders-of-magnitude fewer parameters while attaining similar performance (Li & Liang, 2021; Lester et al., 2021). The parameter space for model h_0 is the set of *soft prompts*, which are matrices $\mathbb{R}^{L \times d}$, where L is a sequence length hyperparameter and is d the dimension of the pre-trained T0 embeddings. Each row of the soft prompt mimics the embedding of a token, but the soft prompt need not correspond to the embedding of any actual token sequence. This matrix is prepended to the input embedding and the output of h_0 is computed with the frozen T0 model. The initial guess at h_0 (i.e., the initial soft prompt vector for use in co-training) is the repeated embedding of the [PAD] token. Since T0 was trained to perform well at zero-shot learning with prompts, this provides a good initial hypothesis. We co-train this model with a pre-trained DeBERTa representation as $\phi_1(X)$ and the last few layers of DeBERTa as h_1 . This is shown in Figure 1, right.

We apply our approach to standard few-shot and zero-shot tasks and find that (i) iteratively co-training models using unlabeled data consistently improves performance, (ii) pseudo-labels from a prompted model are an effective signal for fine-tuning smaller task-specific models, and (iii) this approach can significantly improve results on datasets previously considered difficult for prompt-based learning. We conclude with a brief analysis of success/failure cases and describe high-level criteria required for our method to work.

¹We exclusively use the 3-billion-parameter variant T0-3B.

2. Related work

Prompting and prompt tuning. Lu et al. (2021) find optimal orderings of prompt examples based on an artificially constructed development set. Given the variance in performance across different prompts, others have focused on engineering suitable prompts, manually or otherwise (Liu et al., 2021a). Jiang et al. (2020), Shin et al. (2020), and Gao et al. (2021) use data-driven techniques and language models to automatically generate candidate prompts. Rather than being constrained to human-readable prompts, Li & Liang (2021) and Lester et al. (2021) instead learn a continuous *soft* task-specific “prompt” to condition language models. While effective, these methods typically require nontrivial amounts of labeled data. Several methods try to improve the sample-efficiency of these techniques by using pre-training (Gu et al., 2022) or by combining a hard prompt with a soft prompt (which we also do in this work) and tuning the soft prompt on a few labeled examples (Liu et al., 2021b). In contrast, our technique requires no additional pre-training steps, applies even in the zero-shot setting, and further boosts performance by allowing two different models to learn from each other.

Another line of work uses the outputs from a prompted language model as weak labels, as we do in this work. Wang et al. (2021) propose to train smaller models on labels from GPT-3 to reduce annotation cost, but they train from individual, uncalibrated prompts and do not attempt to refine the prompt model alongside the smaller model. Schick & Schütze (2021) fine-tune a separate RoBERTa model for *each* prompt using a small amount of labeled data. They next aggregate the outputs of these individual fine-tuned

models as a soft pseudo-label and train a final model to match the soft aggregation. In contrast, we train a single BERT-style model on the ensembled prompt output without any additional labeled data. We use this model to refine the ensemble parameters (and vice-versa). In our approach we only use prompt outputs as training signal, and we consider different types of prompts (open-ended instead of cloze).

Self-training for few-shot text classification. Our work relies on access to a large amount of unlabeled data to iteratively grow a confidently-labeled training set for each model. Similarly, *self-training* first trains a model on a small set of initial data, uses the trained model to produce pseudo-labels on a set of unlabeled data, and then iteratively includes the confidently pseudo-labeled data as new training labels (Scudder, 1965). In the context of few-shot text classification, Mukherjee & Awadallah (2020) develop an uncertainty-aware technique for choosing which data points to include, which requires a small amount of labeled data. Karamanolakis et al. (2019; 2021) employ self-training and iterative co-training with weak supervision as the initial label signal, and they similarly use a neural network with pretrained embeddings as a downstream model. However, they explore hand-written or keyword-based rules as weak supervision, in contrast to the present work, where we derive our weak signals from prompted models. The parameterization of h_0 in our partial access setting is similar to the weighting they use to combine rules.

Co-training. Co-training dates back to Blum & Mitchell (1998), who assumed that $\phi_0(X)$ and $\phi_1(X)$ are two distinct views and conditionally independent given the true label. Under this strict condition, they proved that the algorithm finds a good classifier after just one step. Many subsequent analyses (e.g., Dasgupta et al., 2002; Balcan et al., 2005) relax this condition, showing that views can be dependent or even identical as long as certain relationships hold between the models being trained (essentially, they are “different enough”). In a similar vein, Wei et al. (2020) give a theoretical explanation of why (and when) models can learn to be more accurate than the pseudo-labels used to train them. We take implicit advantage of these results in our work. The views we use are highly dependent, and yet the models we train are often able to outperform the pseudo-labels we used to train them in each co-training iteration.

3. Co-training with prompting

The skeleton of our approach is shown in Algorithm 1 (full detail is provided in Algorithms 4 and 5 in the supplement). First, a hypothesis h_0 over view ϕ_0 is initialized such that its initial predictions are reasonable. (We discuss initialization in depth in the following sections.) Next, we obtain the confidently labeled training data L_0^0 , which is a subset of the unlabeled data points, together with pseudo-labels for

Algorithm 1 Co-training algorithm

```

input  $\mathcal{U} = \{x_n\}_{n=1}^U$  unlabeled examples
input  $\{(x_j, y_j)\}_{j=1}^k$  labeled examples (optional)
input initial coverage  $\beta$ , coverage increase  $\beta'$ 
 $h_0 \leftarrow \text{InitClassifier}(\phi_0)$ 
for  $t$  in  $\{0, \dots, T - 1\}$  do
     $\tilde{\beta} \leftarrow \beta + t\beta'$ 
    // GetConfData* defined in Algorithms 2, 3
     $L_0^t \leftarrow \text{GetConfData}^*(\mathcal{U}; h_0, \phi_0, \tilde{\beta})$ 
     $h_1 \leftarrow \text{Train}(\phi_1, L_0^t)$ 
     $L_1^t \leftarrow \text{GetConfData}^*(\mathcal{U}; h_1, \phi_1, \tilde{\beta})$ 
     $h_0 \leftarrow \text{Train}(\phi_0, L_1^t)$ 
end for
return  $(h_0, h_1)$ 

```

those points from h_0 . In iteration t , we select a $\beta + t\beta'$ fraction of the data. (We discuss techniques for selection of confident data in Section 4 and the choice of β and β' in Section 5.) These confidently-labeled points are then used to train a model h_1 on view ϕ_1 , and h_1 's confidently-labeled data is extracted as L_1^0 . This is used to train a new h_0 , and the process continues for T steps. `Train` performs standard supervised training on the pseudo-labels for that iteration. In this section, we give details for how to construct the views ϕ_0 and ϕ_1 , the hypothesis classes we use for the model h_0 , and the initialization schemes for h_0 in both the partial access and full access settings.

3.1. Partial access setting: co-training a label model

In the usual few-shot setting with prompting (also known as *in-context learning*), k labeled examples $\{(x_i, y_i)\}_{i=1}^k$ are converted into a single natural language prompt following a template (example below). We call this prompt k -shot, since it uses k labeled examples. Instead of using one k -shot prompt, in this work we use k one-shot prompts, only including one example in the template at a time. This gives us k outputs. Separating out the signal from each labeled example in this way allows us to combine the examples more effectively than the one k -shot prompt model.

View. Let $\phi_0^{(i)}(x) \in \mathbb{R}^{|V|}$ be the vector of probabilities output by GPT-3 on input x formatted in a one-shot prompt with labeled example (x_i, y_i) . Here $i \in \{1, \dots, k\}$ and V is a subset of the full token vocabulary—the *verbalizer tokens*—and consists of the “label tokens” for the prompt as well as other tokens related to the label. For example, in sentiment analysis, if x_1 is “this movie was great!”, $\phi_0^{(1)}(x)$ is GPT-3’s output on:

```

Review: this movie was great!
Positive or Negative? Positive
Review: {{x.review}}
Positive or Negative?

```

and V might include the label tokens `Positive / Negative` and related tokens such as uncased label tokens

or synonyms.² For example, `Date` appears in the GPT-3 output for TREC question classification even though the closest label token is `Number`. On binary tasks, `yes/no` sometimes appear for prompts whose pre-specified label tokens are `Yes/No`. By including these other tokens in V , we allow the label model to automatically learn the correct associations during co-training.

To select V in a task-agnostic way, we obtain the 10 most likely predictions from GPT-3 on each prompt/example pair. After obtaining these outputs for the entire unlabeled training set, we sort all the observed tokens by the total probability assigned to them across all examples. We select the top 25% of the tokens with nonzero probability mass as V . This ensures that frequently-output tokens appear in the feature set. If a token in V did not appear in the top 10 outputs for a prompt/example pair, we set $\phi_0^{(i)}(x)$ to 0 for that token instead of re-querying the model to get its probability.

By concatenating $\phi_0^{(i)}(x)$ for each of the k labeled examples, we obtain a matrix $\phi_0(x) \in \mathbb{R}^{k \times |V|}$, the first view for co-training. The second view, $\phi_1(x)$, is the frozen representation of a pretrained model like DeBERTa (He et al., 2021). In our experiments, we use the representation in the penultimate layer as $\phi_1(x)$, and the hypothesis class over this view is the last layer and the linear classifier.

Hypothesis class. This leaves the hypothesis class for model h_0 : how do we combine k prompt signals into one pseudo-label? Probabilities from these models are often miscalibrated (Zhao et al., 2021), and thus averaging or majority voting does not yield good results. Instead, we propose to learn a *label model* that scales each prompt vector $\phi_0^{(i)}(x)$ by a prompt-specific *calibration matrix* $W^{(i)}$ before averaging. The combined architecture for this model is given by $h_0(x; W, \alpha)$:³

$$\begin{aligned} \mathbf{l}_i &= \text{ReLU}\left(W^{(i)}\phi_0^{(i)}(x)\right); \\ h_0(x; W, \alpha) &= \text{softmax}\left(\sum_{i=1}^k \alpha_i \mathbf{l}_i\right), \end{aligned} \quad (1)$$

where $\alpha \in \mathbb{R}^k$ is a vector of weights for ensembling the scaled prompt outputs. The $\text{ReLU}(\cdot)$ allows the model to easily ignore particular prompt/label combinations. For example, if prompt j has very poor precision when it outputs label z , setting $W_{zz}^{(j)}$ to be negative causes l_{jz} to be 0. Note that we directly calibrate probabilities rather than *log*-probabilities, following Zhao et al. (2021) (i.e., $\phi_0^{(i)} \in [0, 1]^{|V|}$ and $\|\phi_0^{(i)}\|_1 = 1$). In each iteration of co-training, we train this model using the standard cross-entropy loss on the confident data for that iteration.

²In the running sentiment analysis example we might have $V = \{\text{Negative, Positive, No, Yes, no, yes, bad, good, ...}\}$.

³Here we use W to refer to $\{W^{(i)}\}_{i=1}^k$.

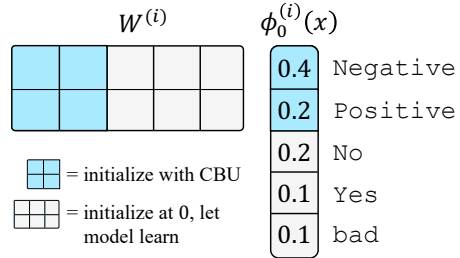


Figure 2. Initialization for the label model parameters $W^{(i)}$ together with example input probability vector $\phi_0^{(i)}(x)$. The weights for the pre-specified label tokens `Positive/Negative`, shown in blue, are initialized using CBU (Zhao et al., 2021). The weights for other select tokens (shown in gray) are initialized to 0, but learned during subsequent co-training steps. This allows the model to automatically learn additional verbalizer tokens during the co-training iterations.

Initialization. If there are l labels for a task, then the calibration matrices $W^{(i)} \in \mathbb{R}^{l \times |V|}$. The $l \times l$ block corresponding to the label tokens is initialized with Calibrate-Before-Use (CBU) (Zhao et al., 2021), and the parameters corresponding to other tokens in V are set to 0. This block structure ensures that only the label tokens (e.g., `Negative`, `Positive`) are used at initialization, but subsequent iterations of co-training can refine the verbalizer with nonzero weights for the other tokens.

CBU first computes the probability vector $\phi_0^{(i)}(x_{cf})$ (restricted to the label tokens only) on content-free inputs x_{cf} (e.g., `N/A` or the empty string), and then uses these as a scaling factor. This ensures the scaled prompt outputs are neutral on truly neutral inputs, improving calibration. Assuming the label tokens are the first elements of V , we set:

$$W^{(i)} = \left[\text{Diag}\left(\frac{\mathbf{1}}{\phi_0^{(i)}(x_{cf})}\right); 0 \right].$$

Figure 2 shows the initialization process for $W^{(i)}$ visually. We also set $\alpha_i = 1$ for each $i \in \{1, \dots, k\}$ to initially weight each prompt equally in the ensemble.

3.2. Full access setting: co-training a soft prompt

In this setting, our prompt model is the T0 model (Sanh et al., 2022), which achieves zero-shot generalization by fine-tuning T5 (Raffel et al., 2020) on multiple tasks whose labeled examples have been transformed into natural language question-answer pairs. Since T0 is publicly available and smaller than GPT-3, we can introspect the model and compute gradients in this case.

View. We set $\phi_0(X)$ to the initial word embeddings of T0 and leave $\phi_1(X)$ and h_1 unchanged (i.e., ϕ_1 is the penultimate layer of a pretrained DeBERTa representation).

Hypothesis class. The model h_0 is parameterized by a continuous *soft prompt* (Li & Liang, 2021; Lester et al., 2021). Concretely, letting $d = 2048$ be the dimension of the T0 word embeddings, a soft prompt is a matrix of parameters $P \in \mathbb{R}^{L \times d}$, where L is a sequence length hyperparameter. Each row of the soft prompt acts like the embedding of a “token” (but needn’t correspond to the embedding of any real token—i.e., there are no constraints on P). The hypothesis $h_0(x; P)$ is thus given by prepending the soft prompt to the input word embedding sequence and using the concatenation ($P; \phi_0(X)$) as input to T0. The subsequent T0 layers are frozen and not updated during training. Given enough labeled data, soft prompt tuning can match the performance of full-fine-tuning with far fewer trainable parameters (Lester et al., 2021; Le Scao & Rush, 2021).⁴

Initialization. T0 is specifically trained to perform well at zero-shot tasks with a variety of hard prompts, so using a hard prompt out-of-the-box gives good initial performance. Hence, to initialize a *soft prompt hypothesis*, we encode the input using a hard prompt and then set the soft prompt to be the repeated embedding of the tokenizer’s padding token. Using the RTE dataset (Dagan et al., 2005) as a running example, we first encode the input using a hard prompt, where each input example x is formatted as:

```
{{x.premise}}
Question: {{x.hypothesis}} True or False?
```

We then set h_0 to be the repeated embedding of the T0 padding token, i.e., at initialization the T0 model sees:

```
[PAD]...[PAD]{{x.premise}}
Question: {{x.hypothesis}} True or False?
```

This combination of hard prompt encoding with soft prompting differs from the usual soft prompting setup (Li & Liang, 2021; Lester et al., 2021). We discuss this issue in more depth in Section B.3.

4. Selecting confident data

The key step in co-training is selecting confidently-labeled data for use in the next training iteration. The literature on co-training has identified a large number of methods for performing this data selection (GetConfData, in Algorithm 1). We consider two simple approaches in this work: *model confidence* and *cut statistic*. In both cases, we specify an initial coverage fraction β and a coverage increase fraction β' . Given U unlabeled examples, the amount of pseudo-labeled data in round $t \geq 0$ is therefore $U(\beta + t\beta')$.

Model confidence. For model confidence, we sort every example by the scores output by each model and select the top $\beta + t\beta'$ fraction in iteration t . While simple, this can result in very imbalanced updates to the pseudo-labeled

dataset if the model is only confident for one label or if one label is inherently more noisy than the others. If additional knowledge regarding the marginal label distribution is available (e.g., approximate label balance or a constraint on the minimum label frequency), we can imbue this knowledge into the data selection process by grouping examples by their predicted label and then performing the sort-and-select procedure for each label separately. Knowledge of the approximate label balance is a standard assumption in weak supervision (e.g., Fu et al., 2020), but we make a much weaker assumption when using the model confidence ranking: we assume we know a *lower bound* γ such that for all labels y , $\mathbb{P}[Y = y] \geq \gamma$. We set $\gamma = 0.01$, i.e., that every class accounts for at least 1% of the data. The detailed procedure for confident data selection using model confidence is shown in Algorithm 2 (supplement).

Cut statistic. The cut statistic is a ranking heuristic that uses the view geometry more than the model confidence approach (Muhlenbach et al., 2004; Zhang & Zhou, 2011). Suppose we want to select data confidently labeled by a model over view $\phi(X)$ (we omit the subscript i for clearer notation). First, we form a graph $G = (V, E)$ with one vertex for each unlabeled training example and edges connecting vertices who are K -nearest neighbors in $\phi(X)$ (or a representation related to $\phi(X)$ —for example, for T0 we can use a contextual representation from inside the model instead of the uncontextual embeddings ϕ_0).

Let $\hat{Y}(X) = \arg\max h(\phi(X))$ be the hard pseudo-label assigned to input X by model h . We say an edge (x_u, x_v) is *cut* if $\hat{Y}(x_u) \neq \hat{Y}(x_v)$. Intuitively, we can feel confident about examples that have few cut edges, since they have the same label as most of their neighbors. Regions of G with high noise are less likely to be correctly labeled. The cut statistic heuristically quantifies this idea to rank examples.

Suppose (as a null hypothesis) that the labels \hat{Y} were sampled i.i.d. from the marginal distribution $\mathbb{P}[\hat{Y} = y]$ (i.e., independently of X). For vertices u and v corresponding to examples x_u, x_v , define $I_{uv} = \mathbb{I}[\hat{Y}(x_u) \neq \hat{Y}(x_v)]$. Consider the test statistic: $J_u = \sum_{v \in N(u)} w_{uv} I_{uv}$, where $w_{uv} = 1/(1 + \|\phi(x_u) - \phi(x_v)\|_2)$ are edge weights that decrease as the distance between u and v increases, and $N(u)$ are the neighbors of u . The mean of J_u under the null hypothesis is: $\mu = (1 - \mathbb{P}[\hat{Y}(x_u)]) \sum_{v \in N(u)} w_{uv}$, and the variance is: $\sigma^2 = \mathbb{P}[\hat{Y}(x_u)](1 - \mathbb{P}[\hat{Y}(x_u)]) \sum_{v \in N(u)} w_{uv}^2$. Following Zhang & Zhou (2011), we approximate the distribution of J with a normal distribution of mean μ and variance σ^2 . Then we can rank examples x_u by the left-sided tail probability for J_u (*lower* is better). If J_u is much smaller than expected, then the total cut edge weight is much smaller than expected under the null hypothesis. To select confident data, we sort examples by J_u and choose the top $\beta + t\beta'$ fraction in iteration t . The detailed procedure for

⁴We use $L = 20$ in our experiments, following Lester et al. (2021), so the soft prompt has $20 \times 2048 = 40960$ parameters.

confident data selection using the cut statistic is shown in Algorithm 3 (supplement).

Relabeling. Pseudo-labels from previous iterations can either be re-used or thrown out. If the initial hypothesis has high precision but low coverage, it is typically preferable to re-use the pseudo-labels from previous iterations, since as coverage increases the quality of the pseudo-labels is likely to go down. On the other hand, if the models being trained are capable of correcting incorrect pseudo-labels, it is preferable to relabel, since this can improve the quality of the training data in each iteration. We exclusively use the latter, since we found that the pseudo-label accuracy on the covered subset of data often *increased* with more iterations. The original co-training algorithm (Blum & Mitchell, 1998) builds L cumulatively, but subsequent co-training procedures also use relabeling (Zhang & Zhou, 2011).

5. Experiments

Datasets. We investigate the benefit of co-training on several standard natural language benchmarks, focusing on datasets with a large gap between the best prompt-based methods and fully-supervised learning (Wang et al., 2019b;a; Brown et al., 2020). We use the **RTE** (Dagan et al., 2005), **CB** (De Marneffe et al., 2019), **TREC** (Voorhees & Tice, 2000), and **BoolQ** (Clark et al., 2019) datasets. Full details for these datasets are in Appendix B. In the partial access setting, we do not evaluate on BoolQ due to the large amount of GPT-3 quota required for labeling. In the full access setting, we do not evaluate on TREC as T0 was pretrained on TREC.

Training methods, partial access. In the few-shot setting we randomly select $k = 4$ training examples from each dataset until the initial label model assigns every pseudo-label at least $\gamma\beta U$ times (i.e., we resample prompts until we can initialize the label model in accordance with the constraint that $\mathbb{P}[Y = y] \geq \gamma$ for all y). While larger k might improve performance, $k = 4$ gives a good balance between performance and GPT-3 quota usage. (Indeed, with our 4 one-shot prompts, we are able to beat the GPT-3 32-shot accuracy on CB).

In each co-training iteration, we train the label model over view ϕ_0 using Adam with learning $1e-4$, weight decay $5e-3$, and batch size 64 for 40 epochs. We fine-tune the last layer and pooler of DeBERTa-large over ϕ_1 for 20 epochs using Adam with learning rate $1e-5$, weight decay 0.01, batch size 16. All parameters were frozen except the last language model layer, the pooler, and the linear classification layer. In order to avoid indirect label leakage, we did not tune these hyperparameters and instead chose common hyperparameters used for these types of models. For early stopping, each model was evaluated every epoch on a *pseudo*-labeled validation set and the best model checkpoint was chosen

based on *balanced accuracy* on the pseudo-labels at the end of each round. Using the balanced accuracy (average of the recall for each label) avoids collapsing to the majority class even when the pseudo-labels are relatively imbalanced. This validation set was sampled uniformly from the training set to give a training/validation split of 90%/10%.

To determine β, β', T for co-training, we performed a light hyperparameter search based on performance on a *gold*-labeled validation set of 500 examples sampled from the TREC training set.⁵ This resulted in the the following values: initial coverage of $\beta = 0.5$, per-step coverage increase of $\beta' = 0.1$, and total co-training steps $T = 5$. We emphasize that this gold validation set was *not* used during any co-training iteration (e.g. for model selection, early stopping, learning rate tuning, etc.) We set the minimum label frequency $\gamma = 0.01$ and did not tune this value. We used model confidence to add confident data in view 0 and the cut statistic to add confident data in view 1.⁶ We used the [CLS] token embedding in the last layer of DeBERTa for cut statistic nearest neighbors in view 1. We used $K = 20$ nearest neighbors for the cut statistic and performed no tuning on this value. Our code is publicly available⁷.

Training methods, full access. In the zero-shot setting, for training the soft prompt over view $\phi_0(x)$ we mainly used the hyperparameters suggested by Lester et al. (2021), which were obtained by performing gold soft prompt tuning using T5 on SuperGLUE. We used Adafactor with constant learning rate 0.3, weight decay $1e-5$, and batch size 24 for 30000 training steps. For DeBERTa-large, we used the same hyperparameters as in the partial access setting. As in the partial access setting, we used balanced pseudo-label accuracy to select the best model checkpoint at the end of each training round. We used the cut statistic for confident selection in *both* views, since with T0 we have access to the internal embeddings, unlike with GPT-3. We used the T0 decoder’s contextual embedding for the first decoded token to compute nearest neighbors for the view 0 cut statistic. Training details (e.g., β, β', T , etc.) are otherwise exactly the same as in the partial access setting.

During training, the pseudo-label for each example is first mapped to a token that matches the hard prompt (e.g. $0 \rightarrow \text{True}$ and $1 \rightarrow \text{False}$ for the RTE example above). These token labels are then mapped to embedding indices using the T0 tokenizer, and the soft prompt is trained via

⁵Hence our few-shot experiments on TREC are not few-shot in the truest sense of the term.

⁶This works better than using the cut statistic in both views—since the cut statistic relies heavily on good nearest neighbors, it makes the most sense in a view that already has a good distance function for examples (the pretrained DeBERTa representation).

⁷<https://github.com/clinicalml/cotrain-prompting>

regular sequence-to-sequence training with the maximum likelihood objective. This is identical to the soft prompt training technique from Lester et al. (2021).

Caveat. As noted by Perez et al. (2021), much current work on prompt-based learning does not constitute “true” few-shot/zero-shot learning as they often implicitly assume access to a small labeled set to select various model configurations (e.g., prompts and hyperparameters). Insofar as we inherit such configurations from existing work, our work is similarly not few-shot/zero-shot in the strictest sense, although we tried to minimize such issues by using exactly the same co-training parameters (β, β', T, γ) and model hyperparameters for all datasets. (We also did not perform an extensive tuning of these parameters.) While we are encouraged by the observation that model configurations seem to work well across diverse datasets, investigating co-training in the context of true few-shot/zero-shot learning (Schick & Schütze, 2021) is an important avenue for future work.

Baselines. For baselines we compare against:

- *GPT-3 32-shot*: From Brown et al. (2020). 32 examples combined in one prompt. Uncalibrated.
- *Calibrate Before Use*: Performance of CBU using 4-shot prompts (from Zhao et al. (2021)).
- *Prompt-based FT*: Our reproduction of the fine-tuning method from Gao et al. (2021), using 2 labels per class.
- *Snorkel on GPT-3 output*: Snorkel generative label model (Ratner et al., 2016; 2020), which aggregates over the four GPT-3 1-shot outputs without using any labeled data.
- *Snorkel + DeBERTa-large*: DeBERTa-large fine-tuned on outputs from Snorkel label model using the same hyperparameters as our co-training methods. This baseline uses all of the pseudolabeled data at once for one round of training, whereas co-training performs multiple rounds.
- *Label Model (no co-training)*: the label model after initialization with (1).
- *Label Model (self-training)*: identical to co-training, but we use the label model for both views (i.e., we train it on its own confident outputs). We use the same hyperparameters that we use for our co-training method.

The baselines that use (roughly) the same amount of labeled data as our method are shown in the top section of Table 1 and Table 2. The bottom sections contain baselines that use more labeled data, including oracle upper-bounds based on full training data. Training details for the baselines are in Section B.2.

5.1. Results

Table 1 shows the results for the partial access setting, where we co-train the label model, which calibrates and combines

multiple GPT-3 outputs, with a smaller pretrained model (DeBERTa-large). For view 0, our co-trained label model (*Label Model + co-training*) improves over the initial label model (*Label Model before co-training*) and the average performance of GPT-3 4-shot before (*GPT-3 4-shot*) and after (*Calibrate Before Use*) calibration. We also improve over *Snorkel on GPT-3*, which, like our method, uses unlabeled data to combine the outputs of our four 1-shot prompts. Finally, co-training outperforms self-training the label model on its own in all cases, indicating that using two complementary models improves performance.

For CB, the co-trained label model outperforms *GPT-3 32-shot* despite only using 4 labeled examples. This suggests that using unlabeled data to learn to ensemble k 1-shot prompts can be more label-efficient than putting all k labeled examples in one prompt. For TREC and CB, the co-trained label model also outperforms prompt-based fine-tuning (*Prompt-based FT* (Gao et al., 2021)) with the same amount of labeled data (Prompt-based FT also uses a gold-labeled validation set of k examples per class, whereas our method only uses a pseudo-labeled validation set). For RTE and CB, we nearly match the fully-supervised performance on view 0 (*Label Model on full train*), suggesting that co-training is able to extract nearly all of the signal from the GPT-3 probabilities in these cases without using any extra labeled data.

For view 1, the co-trained DeBERTa-large model outperforms all of the baselines that use the same amount of label information. For RTE and TREC, it outperforms Prompt-based FT even when the latter uses 4x (for RTE) and 12x (for TREC) the number of labeled examples (*Prompt-based FT with 8 labels per class*). This suggests that the pseudo-labels provided by (a learned ensemble of) prompts are an effective training signal for smaller models. The co-trained DeBERTa-large also outperforms DeBERTa-large trained with the Snorkel label model. This baseline uses all of the pseudolabeled data for one round of training, whereas our method gradually increases the training set size across multiple iterations. The disparity between the methods shows that multiple careful iterations are key for good performance.

Table 2 shows the results for the full access setting with T0. For RTE and CB, co-training improves on the performance of the initial zero-shot prompt (*T0-3B zero-shot (no co-training)*). For RTE, the co-trained view 0 and view 1 models nearly match the performance of their fully-supervised counterparts. The difference in co-training performance on RTE in Table 1 and Table 2 shows the benefit of having full access to h_0 . Since we can introspect the prompt model, we can use the cut statistic in both views. In the first step of co-training, the cut statistic on view 0 selects confident data with 90%-accurate pseudo-labels. The confident data selection on CB is similarly good: in the first co-training

Table 1. Few-shot learning results. (Top) Results against various baselines that use exactly 4 labels per dataset (except for Prompt-based FT, which uses 8 labels for CB and 12 labels for TREC, since this approach uses labels at the “per-class” level). GPT-3 and CBU results are copied from Zhao et al. (2021), while we train our own Prompt-based FT (Gao et al., 2021) and Snorkel (Ratner et al., 2020) models. (Bottom) Results from baselines trained on more data (for reference only). Standard deviation (when applicable) numbers are given by 4 runs over prompts (GPT-3, CBU) or random seeds (Snorkel, Prompt-based FT, Co-training). † rows show accuracy on the private SuperGLUE test set. Otherwise, the accuracies are on the public SuperGLUE validation sets, which we treated as a test set.

Model	View	RTE (2-class)	CB (3-class)	TREC (6-class)
GPT-3 4-shot (from Zhao et al. (2021))	*	58.7 (11.9)	45.2 (19.4)	60.2 (7.6)
Calibrate Before Use (CBU) (Zhao et al., 2021)	*	60.4 (8.1)	60.7 (6.7)	69.7 (1.4)
Prompt-based FT (Gao et al., 2021)	*	52.8 (0.9)	84.4 (3.2)	54.8 (2.9)
Snorkel on GPT-3 (Ratner et al., 2020)	ϕ_0	59.6 (0.0)	70.2 (0.8)	65.2 (0.0)
Snorkel on GPT-3 + DeBERTa-large	ϕ_1	67.2 (0.5)	81.6 (2.2)	63.3 (0.4)
Label Model (no co-training)	ϕ_0	62.8	76.8	77.2
Label Model + self-training	ϕ_0	53.5 (0.6)	45.7 (5.7)	75.4 (0.2)
Label Model + co-training	ϕ_0	64.9 (1.1)	83.5 (2.3)	78.3 (1.2)
DeBERTa-large + co-training	ϕ_1	67.4 (2.3)	86.2 (3.2)	80.6 (1.1)
Label Model on full train	ϕ_0	67.8 (0.5)	82.7 (0.8)	91.9 (1.1)
DeBERTa-large on full train	ϕ_1	93.3	95.2	96.7
GPT-3 32-shot† (Brown et al., 2020)	*	69.0	75.6	*
Prompt-based FT with 4 labels per class	*	48.3 (3.3)	84.4 (4.4)	58.8 (5.9)
Prompt-based FT with 8 labels per class	*	56.1 (3.2)	87.5 (0.0)	80.0 (4.8)

Table 2. Zero-shot learning results with T0 as the initial view 0 model and DeBERTa as the second model. We also show the results trained on the full dataset in the bottom two rows. For T0-3B (best), we take the *best-performing* prompts from Sanh et al. (2022) and replicate their results as exact numbers for each prompt were not provided in the original paper. T0-3B zero-shot (no co-training) reports the initial accuracy of the prompt used by our method (see Appendix D for the full list of prompts used by our algorithm, which we selected before looking at their test performance). Standard deviations are not provided in this case as even a single run takes a nontrivial amount of time.

Model/Algorithm	View	RTE	CB	BoolQ
T0-3B (best) (Sanh et al., 2022)	ϕ_0	68.9	66.1	59.1
T0-3B zero-shot (no co-training)	ϕ_0	68.9	58.9	56.4
T0-3B soft prompt + co-training	ϕ_0	87.0	67.9	49.1
DeBERTa-large + co-training	ϕ_1	86.3	67.9	48.9
T0-3B soft prompt on full train	ϕ_0	90.6	80.4	86.9
DeBERTa-large on full train	ϕ_1	93.3	95.2	86.1

step, the cut statistic selects pseudo-labels with 89% accuracy. The pseudo-labels extracted by the view 1 model after the first step of co-training (L_1^0) are 98% accurate, so after training on the initial pseudo-labels, the view 1 model is able to select a very high quality training set at coverage $\beta = 0.5$. However, the CB performance is *worse* than in Table 1 despite the strong initial signal in L_0^0 and the near-perfect training data in L_1^0 . Similarly, for BoolQ, co-training makes the soft prompt *worse* than the initial zero-shot model. We explore the reasons behind this below.

When (and how) does co-training work? Figure 3 (left) shows the evolution of the test accuracy for h_0 and h_1 over co-training iterations on TREC (from Table 1). Multiple rounds of co-training increase performance for both views as the models become more precise and the coverage increases. Figure 3 (right) shows the precision of the *confident* data extracted by h_0 for each iteration of co-training, broken down by label. This figure shows two distinct phenomena. For most labels, precision decreases as coverage goes up, as we

expect from usual co-training theory (see e.g. Balcan et al., 2005). However, for label 1, precision actually *increases* over iterations. Model h_1 (DeBERTa) is able to select new confident data for label 1 that is *better* than the weak labels used to train it, which improves the h_0 precision for label 1 in subsequent iterations. For example, in iteration 2, h_0 ’s confident precision for label 1 is 0.39, but after h_1 is trained on that data, it proposes new confident data for label 1 with precision 0.58 (not shown in Figure 3). Pseudo-label correction is one of the benefits of having two complementary models (though it can also happen with a single model with appropriate regularization (Wei et al., 2020)).

Figure 4 shows what can happen when h_0 and h_1 are not complementary enough. The left display shows the accuracy of each model over co-training iterations. The right display shows the balanced accuracy of the confident data extracted from each model. In the first co-training step, h_1 greatly improves over the initial h_0 , and selects extremely accurate confident data (nearly 100% accurate) at coverage




Figure 3. Partial access setting, TREC. Left: Test accuracy vs co-training iteration for the label model h_0 and the DeBERTa model h_1 . Right: precision per label vs co-training iteration, h_0 .




Figure 4. Full access setting, CB. Left: Test accuracy vs co-training iteration for T0-3B (h_0) and the DeBERTa model (h_1). Right: balanced accuracy of confident pseudo-labels extracted from T0-3B (h_0) and DeBERTa (h_1).

$\beta = 0.5$. This improves the performance of the soft prompt for the next iteration (h_0 , left, iteration 1), but the confident balanced accuracy of h_0 sharply decreases. Inspecting the training of h_0 on L_1^0 , the soft prompting procedure appears to have overfit to the pseudo-labels on L_1^0 . Due to the small size of CB (250 training examples), coverage $\beta = 0.5$ and our 90/10 train/val split gives only 112 data points for training h_0 in the first iteration, but the soft prompt is a very flexible hypothesis class. This overfitting causes the accuracy of confident data to decrease, which in turn degrades the performance of h_0 , and eventually the two models converge to almost identical predictors. This suggests that CB does not have enough unlabeled data for co-training to perform well with T0, at least when $\beta = 0.5$. Using larger initial coverage (e.g. $\beta = 1.0$, $\beta' = 0$) or a less flexible hypothesis class for h_0 might improve performance.

Finally, Table 2 showed that co-training *decreased* performance on BoolQ even though the initial soft prompt seemed to have reasonably strong signal (56.4% accuracy). However, the finer-grained statistics are less promising. Let the precision $\text{prec}(\hat{y})$ for a pseudolabel \hat{y} be given by:

$$\text{prec}(\hat{y}) := \mathbb{P}[Y = \hat{y} | \hat{Y} = \hat{y}],$$

and define the ‘‘total noise’’ η as

$$\eta := \sum_y \mathbb{P}[\hat{Y} \neq y | Y = y].$$

Bayes’ rule implies the total noise and the per-pseudolabel precisions are related by

$$\eta := \sum_y \left(1 - \text{prec}(y) \frac{\mathbb{P}[\hat{Y} = y]}{\mathbb{P}[Y = y]} \right).$$

Even under ideal conditions on ϕ_0 and ϕ_1 , $\eta < 1$ is required for learning to work (for binary Y), and the sample complexity depends on $1/(1 - \eta)$ (Blum & Mitchell, 1998). We can use this theory to analyze why co-training fails for BoolQ.

After the first training iteration, with $\beta = 0.5$, we have 5118 training examples in the confident subset. h_1 has $\text{prec}(0) = 0.4$, $\text{prec}(1) = 0.66$, $\mathbb{P}[\hat{Y} = 0] = 0.64$, and $\mathbb{P}[\hat{Y} = 1] = 0.36$. On this subset, the true label balance is $\mathbb{P}[Y = 0] = 0.37$, and $\mathbb{P}[Y = 1] = 0.63$. Plugging these values into the formula above, we find the ‘‘total noise’’ η in the pseudo-labels is 0.93. At the beginning of iteration $t = 1$, when $\tilde{\beta} = 0.6$, the total noise in the confident data assigned by h_0 is even worse, at 0.98. For comparison, the total noise for the initial h_0 on CB is 0.21.

Unfortunately, the same issue persists on BoolQ for different β values. The negative result on BoolQ suggests that the initialization for h_0 needs to have less total noise. A different prompt or a better initial hypothesis (e.g., full T0 instead of T0-3B) could be more amenable to co-training.

6. Conclusion

Our results indicate that using unlabeled data to co-train a prompted model with a smaller model can boost the performance of prompt-based learning on few-shot and zero-shot classification tasks. As a side effect, this procedure also produces a smaller performant model on the task of interest, distilling and refining the knowledge in the large prompted model. Using two complementary models and views allows the models to learn from each other despite training on partially incorrect pseudo-labels. We showed that the benefit of co-training is limited when the initial signal provided by the prompted model is too noisy (BoolQ, full access), when there is not enough unlabeled data to obtain good (pseudo-label) generalization performance (CB, full access), and when there is a large gap in fully-supervised accuracy on view 0 and view 1 (RTE, partial vs full access). Developing methods to overcome these limitations in the context of prompting is an interesting direction for future work.

Acknowledgments

DS and HL were partially supported by NSF AiTF award CCF-1723344. MA was supported by the Takeda Fellowship. Thanks to Dr. Steven Horng of Beth Israel Deaconess Medical Center for providing access to an NVIDIA DGX machine (Horng, 2022), and thanks to NVIDIA Corporation for their donation of two NVIDIA A100 GPUs. Thanks to OpenAI and AI21 for providing quota to access their davinci and Jurassic-Jumbo models (respectively). Finally, thanks to Rebecca Boiarsky and Catherine Wong for their feedback on drafts of this paper and to Aravindan Vijayaraghavan for helpful discussions on co-training theory.

References

- Balcan, M.-F., Blum, A., and Yang, K. Co-training and expansion: Towards bridging theory and practice. *Advances in neural information processing systems*, 17:89–96, 2005.
- Blum, A. and Mitchell, T. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pp. 92–100, 1998.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. BoolQ: Exploring the surprising difficulty of natural yes/no questions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2924–2936, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1300. URL <https://aclanthology.org/N19-1300>.
- Dagan, I., Glickman, O., and Magnini, B. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pp. 177–190. Springer, 2005.
- Dasgupta, S., Littman, M. L., and McAllester, D. Pac generalization bounds for co-training. *Advances in neural information processing systems*, 1:375–382, 2002.
- De Marneffe, M.-C., Simons, M., and Tonhauser, J. The commitmentbank: Investigating projection in naturally occurring discourse. In *proceedings of Sinn und Bedeutung*, volume 23, pp. 107–124, 2019.
- Fu, D., Chen, M., Sala, F., Hooper, S., Fatahalian, K., and Ré, C. Fast and three-rious: Speeding up weak supervision with triplet methods. In *International Conference on Machine Learning*, pp. 3280–3291. PMLR, 2020.
- Gao, T., Fisch, A., and Chen, D. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 3816–3830. Association for Computational Linguistics, August 2021. URL <https://aclanthology.org/2021.acl-long.295>.
- Gu, Y., Han, X., Liu, Z., and Huang, M. Ppt: Pre-trained prompt tuning for few-shot learning. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8410–8423, 2022.
- He, P., Liu, X., Gao, J., and Chen, W. DeBERTa: Decoding-enhanced bert with disentangled attention. In *Proceedings of ICLR*, 2021.
- Hornig, S. Machine learning core. Feb 2022. doi: 10.6084/m9.figshare.19104917.v1. URL https://figshare.com/articles/preprint/Machine_Learning_Core/19104917/1.
- Jiang, Z., Xu, F. F., Araki, J., and Neubig, G. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020.
- Karamanolakis, G., Hsu, D., and Gravano, L. Leveraging just a few keywords for fine-grained aspect detection through weakly supervised co-training. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4611–4621, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1468. URL <https://aclanthology.org/D19-1468>.
- Karamanolakis, G., Mukherjee, S., Zheng, G., and Hassan, A. Self-training with weak supervision. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 845–863, 2021.
- Kumar, S. and Talukdar, P. Reordering examples helps during priming-based few-shot learning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, Online, August 2021. Association for Computational Linguistics.
- Le Scao, T. and Rush, A. How many data points is a prompt worth? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2627–2636, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.208. URL <https://aclanthology.org/2021.naacl-main.208>.
- Lester, B., Al-Rfou, R., and Constant, N. The power of scale for parameter-efficient prompt tuning. In *Proceedings of*

- the 2021 Conference on Empirical Methods in Natural Language Processing, pp. 3045–3059, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.emnlp-main.243>.
- Li, X. L. and Liang, P. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 4582–4597, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.353. URL <https://aclanthology.org/2021.acl-long.353>.
- Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., and Neubig, G. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021a.
- Liu, X., Zheng, Y., Du, Z., Ding, M., Qian, Y., Yang, Z., and Tang, J. Gpt understands, too. *arXiv:2103.10385*, 2021b.
- Lu, Y., Bartolo, M., Moore, A., Riedel, S., and Stenetorp, P. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*, 2021.
- Muhlenbach, F., Lallich, S., and Zighed, D. A. Identifying and handling mislabelled instances. *Journal of Intelligent Information Systems*, 22(1):89–109, 2004.
- Mukherjee, S. and Awadallah, A. Uncertainty-aware self-training for few-shot text classification. *Advances in Neural Information Processing Systems*, 33, 2020.
- Perez, E., Kiela, D., and Cho, K. True few-shot learning with language models. *arXiv preprint arXiv:2105.11447*, 2021.
- Pilehvar, M. T. and Camacho-Collados, J. Wic: the word-in-context dataset for evaluating context-sensitive meaning representations. *arXiv preprint arXiv:1808.09121*, 2018.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21: 1–67, 2020.
- Ratner, A., Bach, S. H., Ehrenberg, H., Fries, J., Wu, S., and Ré, C. Snorkel: Rapid training data creation with weak supervision. *The VLDB Journal*, 29(2):709–730, 2020.
- Ratner, A. J., De Sa, C. M., Wu, S., Selsam, D., and Ré, C. Data programming: Creating large training sets, quickly. *Advances in neural information processing systems*, 29: 3567–3575, 2016.
- Sanh, V., Webson, A., Raffel, C., Bach, S., Sutawika, L., Alyafeai, Z., Chaffin, A., Stiegler, A., Le Scao, T., Raja, A., et al. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations*, 2022.
- Schick, T. and Schütze, H. Exploiting cloze-questions for few-shot text classification and natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 255–269, 2021.
- Schick, T. and Schütze, H. True few-shot learning with prompts – a real-world perspective. *Computing Research Repository*, arXiv:2111.13440, 2021. URL <http://arxiv.org/abs/2001.07676>.
- Scudder, H. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371, 1965.
- Shin, T., Razeghi, Y., Logan IV, R. L., Wallace, E., and Singh, S. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online, November 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.emnlp-main.346>.
- Voorhees, E. M. and Tice, D. M. Building a question answering test collection. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 200–207, 2000.
- Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. Super-glue: a stickier benchmark for general-purpose language understanding systems. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 3266–3280, 2019a.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019*, 2019b.
- Wang, S., Liu, Y., Xu, Y., Zhu, C., and Zeng, M. Want to reduce labeling cost? GPT-3 can help. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pp. 4195–4205, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.findings-emnlp.354>.

- Wei, C., Shen, K., Chen, Y., and Ma, T. Theoretical analysis of self-training with deep networks on unlabeled data. In *International Conference on Learning Representations*, 2020.
- Wei, J., Bosma, M., Zhao, V. Y., Guu, K., Yu, A. W., Lester, B., Du, N., Dai, A. M., and Le, Q. V. Finetuned language models are zero-shot learners. *arXiv:2109.01652*, 2021.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- Zhang, M.-L. and Zhou, Z.-H. Cotrade: Confident co-training with data editing. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 41(6):1612–1626, 2011.
- Zhao, Z., Wallace, E., Feng, S., Klein, D., and Singh, S. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pp. 12697–12706, 2021.

A. Algorithm details

A.1. Relabeling

Pseudolabels from previous iterations can either be re-used or thrown out. If the initial hypothesis has high precision but low coverage, it is typically preferable to re-use the pseudolabels from previous iterations, since as coverage increases the quality of the pseudolabels is likely to go down. On the other hand, if the models being trained are capable of correcting incorrect pseudolabels, it is preferable to relabel, since this can improve the quality of the training data in each iteration.

In iteration t of co-training, we extract a pseudolabeled dataset L_i^t from model h_i and use it to train model h_{1-i} . Let $\mathcal{N} \subset [U]$ be a set of indices that correspond to the data points confidently pseudolabeled by model h_i in this iteration (according to either the model confidence or cut statistic rankings). Define $S = \{(x_n, \hat{y}_n) : n \in \mathcal{N}\}$ as the set of these points together with their pseudolabels $\hat{y}_n := \operatorname{argmax} h_i(\phi_i(x_n))$. Let $L_i^{t-1} = \{(x_n, \tilde{y}_n)\}$ be the confident data used in the *previous* iteration to train model h_{1-i} . For x_n that appear in L_i^{t-1} but where $n \in \mathcal{N}$, we have a choice to make: do we use the old pseudolabel \tilde{y}_n , or the new one \hat{y}_n ? These need not agree, since h_i has been updated.

Let $S' = \{(x_n, \hat{y}_n) \in S \mid \neg \exists y : (x_n, y) \in L_i^{t-1}\}$ be the set of *newly* pseudolabeled examples—points that do not appear in L_i^{t-1} with any pseudolabel. If we choose to re-use the pseudolabels from the previous iteration, the `GetConfData` update is:

$$L_i^t \leftarrow L_i^{t-1} \cup S'$$

On the other hand, if we throw out the previously pseudolabeled data, the update is simply:

$$L_i^t \leftarrow S$$

We exclusively use the latter, since our models can learn to correct bad initial labels (see Figure 3, Label 1). We found that the pseudolabel accuracy on the covered subset of data often *increased* with more iterations. This relabeling technique is different from the original cotraining algorithm (Blum & Mitchell, 1998), which builds L cumulatively, but subsequent cotraining procedures also use relabeling (Zhang & Zhou, 2011).

A.2. Warm starting

Instead of warm-starting the models h_0 and h_1 (initializing them from the output of the previous iteration), we initialize them from scratch each co-training iteration to reduce the effect of a “bad” training iteration and so that we can use the same training hyperparameters for every iteration. This takes advantage of the fact that there exists a robust set of initial hyperparameters that have been shown to work well

for fine-tuning large language models. However, further exploration of warm-starting is an interesting direction for future work, since it may yield significant reduction in the computational burden of co-training.

A.3. Confident data selection

Algorithm 2 shows how to select confident data using model confidence, and Algorithm 3 shows how to select confident data using the cut statistic. As mentioned in the main text, ϕ_0 and ϕ_1 themselves needn’t be the representations used to compute nearest neighbors for the cut statistic. For example, $\phi_0(x)$ for T0 is the non-contextual T0 input embedding of x . Instead of computing nearest neighbors in this view, we use the *contextual* embedding from much later in the T0 model: the final decoder embedding of the first decoded token. This is a function of both $\phi_0(x)$ and the current hypothesis h_0 . Because the embeddings are contextual, this representation has better nearest neighbors than $\phi_0(x)$; because it also takes h_0 into account, these neighbors are adapted to the current task. Similarly, for DeBERTa-large in view 1, we use the [CLS] token embedding in the *last* layer of the DeBERTa representation rather than the penultimate layer, since this layer has been adapted to the task of interest by the time we select confident data.

Validation dataset We use a pseudolabeled validation set to perform model selection during the co-training iterations. Since the confident-data-selection methods can pick out the most precisely pseudolabeled examples (w.r.t. the true label), we also use them to select a *confident* validation set from the larger val set for each `Train` step. In particular, when training model h_i , we use model h_{1-i} to select a $\tilde{\beta} = \beta + t\beta'$ confident fraction of full validation data in each step (the same fraction used for the confident training set). This allows us to use a more precise validation set for model selection.

A.4. Full algorithms

The detailed algorithms for co-training in the partial access setting and full access setting are shown in Algorithms 4 and 5, respectively. Algorithm 4 uses model confidence for view 0 and cut statistic for view 1. Algorithm 5 uses cut statistic for both views. The detailed procedures for model confidence and the cut statistic are shown in Algorithms 2 and 3, respectively. As mentioned in the previous section, the view 1 cut statistic uses the [CLS] token embedding in the last layer of h_1 . In the full access case, the view 0 cut statistic uses the T0 decoder’s hidden state for the first decoded token.

B. Training and dataset details

B.1. Datasets

- **RTE** (Dagan et al., 2005): Binary textual entailment, 2490 training examples, 277 validation examples (our test set). $\mathbb{P}[Y] = (0.5, 0.5)$
- **CB** (De Marneffe et al., 2019): Ternary textual entailment, 250 training examples, 56 validation examples (our test set). $\mathbb{P}[Y] = (0.41, 0.5, 0.09)$
- **WiC** (Pilehvar & Camacho-Collados, 2018): Binary word sense disambiguation, 5428 training examples, 638 validation examples (our test set). $\mathbb{P}[Y] = (0.5, 0.5)$
- **TREC** (Voorhees & Tice, 2000): 6-way question classification, 5452 training examples, 500 test examples. Label balance:
 $\mathbb{P}[Y] = (0.2131, 0.2293, 0.0158, 0.2243, 0.1643, 0.1532)$
- **BoolQ** (Clark et al., 2019): Binary reading comprehension, 9427 training examples, 3270 validation examples (our test). $\mathbb{P}[Y] = (0.38, 0.62)$

B.2. Training details

Prompt-based FT. We fine-tuned the MLM-pretrained RoBERTa-large model using Adam for 1000 steps with batch size 16, learning rate $1e-5$ and weight decay 0.01. We sampled a validation set the same size as the training set while ensuring that the validation set also had an equal number of examples per class. This small validation set was used to select the best model checkpoint in each run and the test results were averaged over four random seeds. This is similar to the “no \mathcal{D}_{dev} ” setting in Gao et al. (2021) in that we didn’t use the small validation set for hyperparameter tuning—we used the same hyperparameters as the “no \mathcal{D}_{dev} ” setting. However, we still allow the method to use the labeled validation set for model selection. We used the same prompt templates as Gao et al. (2021). For RTE, the label words were Yes, No. For CB, the label words were Yes, No, Maybe. For TREC, the label words were Description, Entity, Abbreviation, Person, Number, Location.

Calibrate Before Use (CBU). For x_{cf} , we followed Zhao et al. (2021) and used “N/A”, the empty string, and “[MASK]”. We obtained the GPT-3 outputs for each of these x_{cf} ’s, renormalized the outputs over the label tokens, averaged the re-normalized outputs across the three x_{cf} ’s, and used the *average* result as the scaling factor for $W^{(i)}$. This is identical to Zhao et al. (2021).

Co-training. Following RoBERTa and DeBERTa, we used an MNLI-pretrained checkpoint for RTE and CB (microsoft/deberta-large-mnli on HuggingFaceHub). Otherwise, we used DeBERTa-large

(microsoft/deberta-large). We did not experiment with DeBERTa V2 or V3.

B.3. Soft prompt encoding

As detailed in Section 3, we combine hard prompt encoding with soft prompting. That is, we format the input using a hard prompt, and combine this *formatted* input embedding with the soft prompt matrix. This differs from the usual soft prompting setup (Li & Liang, 2021; Lester et al., 2021), where the input is encoded more neutrally, without a natural-language hard prompt:

```

sentence1:  {{x.premise}}
sentence2:  {{x.hypothesis}}

```

A priori, this difference in input encoding could affect the performance of soft prompt tuning and the zero-shot performance of the initial prompted model. However, the full-training-dataset soft-prompt tuning baseline in Table 2 (*T0 soft prompts on full training set*) uses our hard prompt encoding + soft prompting, and it matches fully fine-tuned DeBERTa-large. This suggests that the accuracy loss from choosing a hard prompt (at least for the prompts that we chose) is minimal.

Using the hard prompt encoding might improve the label efficiency of soft prompt tuning, since the soft prompt parameters can focus on “fixing up” the given hard prompt instead of learning a prompt-like embedding from scratch. On the other hand, if the hard prompt performs poorly, the hard prompt encoding might put an unnecessary upper limit on the soft prompt tuning performance, since the soft prompt may not be able to “undo” the hard prompt performance. Liu et al. (2021b) also combine hard and soft prompts in a similar manner. An in-depth comparison between the neutral encoding from the traditional soft-prompting setup and the hard prompt + soft prompt encoding we propose is an interesting direction for future work.

B.4. Hardware

All models were trained on two NVIDIA A100 80Gb GPUs using PyTorch and the Transformers library (Wolf et al., 2020). For the partial access setting, a full run of $T = 5$ co-training iterations with DeBERTa-large takes roughly two hours on this hardware. For the full access setting, a full run of $T = 5$ co-training iterations with T0-3B and DeBERTa-large takes roughly 40 hours. While soft-prompt-tuning is more parameter-efficient than full-fine-tuning, it still requires backpropagation through the entire model and is therefore fairly compute- and memory-intensive. Additionally, we found that full float32 performed much better than float16. Hence, we used float32, which further impacted memory usage and computation time. BF16 may give (nearly) the best of both worlds, but we did not extensively test with this format.

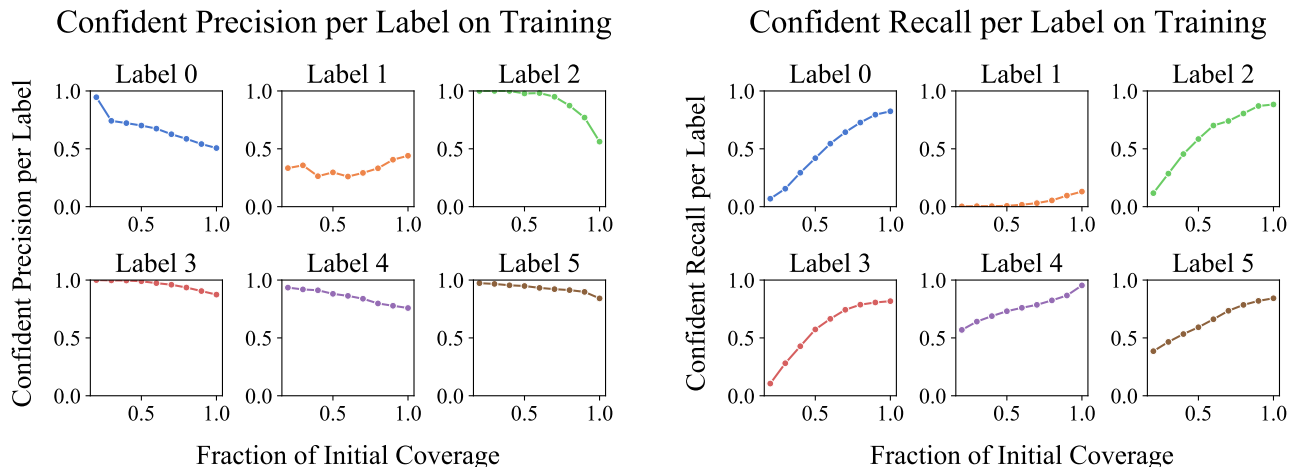


Figure 5. Precision (left) and Recall (right) versus β for each label in the initial confident set L_0^0 , extracted from the initial label model using the model confidence method (Algorithm 2). The precision of some labels (e.g., 2, 3) begins to decline more sharply after $\beta = 0.5$. This gives additional evidence for our choice of $\beta = 0.5$: it trades off between the initial precision of L_0^0 and the coverage for each label. At smaller values of β , there are no pseudolabeled examples for label 1 and very few for label 0. At larger values of β , the precision of the other labels is worse.

C. Additional co-training analysis

In this section, we provide more information regarding the evolution of h_0 and h_1 over the co-training iterations for TREC. We focus on the TREC dataset since its 6 classes enable us to investigate more complex co-training dynamics.

To see the effect of β on the quality of the initial confident data L_0^0 , we plot the precision and recall for each label for different values of β in Figure 5. This figure indicates that the tradeoff when choosing β is between having high precision for each label (lower β) and having enough coverage for each label to train on (high β).

To show how co-training affects label balance across multiple iterations, we plot the total variation distance between the true label balance and the balance estimated using the pseudolabels in each iteration’s confident data L_0^t . Figure 6 indicates that this distance decreases with co-training iterations, so the label model automatically learns to have a balance closer to the unseen true balance.

In Figures 7, 8, and 9, we plot the recall, *normalized coverage*, and precision for each label in L_0^t and L_1^t . The normalized coverage for label j is the number of examples with pseudolabel j divided by the number of examples with true label j ; it separates coverage from the precision, unlike recall. By comparing the evolution of label curves in Figure 8 and 9, we can see that the models tend to add more confident data when they are more precise and add less confident data when they are less precise, which is the desired behavior. Additionally, these figures show two different ways in which co-training works to improve models: “coverage-expansion” and “pseudolabel-correction.” In the coverage-expansion regime, the precision for a label slightly

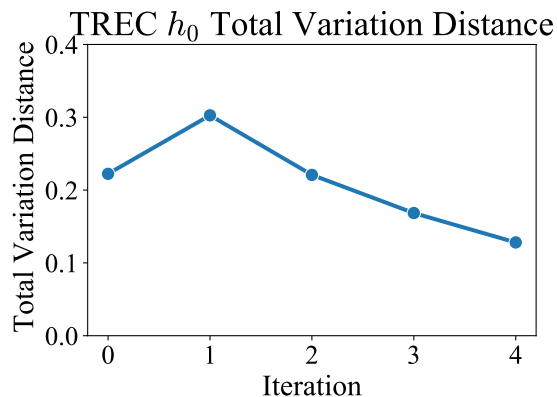


Figure 6. Total variation distance between the true label balance and the label balance estimated from the pseudolabels L_0^t at each iteration. As co-training iterations proceed, the label model automatically learns a balance closer to the true unseen label balance. decreases as iterations increase, but the coverage improves; this regime was predicted by early work on co-training (Balcan et al., 2005). In the pseudolabel-correction regime, both precision and coverage increase, because models are able to learn to be more accurate than the pseudolabels used to train them. Wei et al. (2020) give a theoretical explanation of this in the context of self-training, rather than co-training.

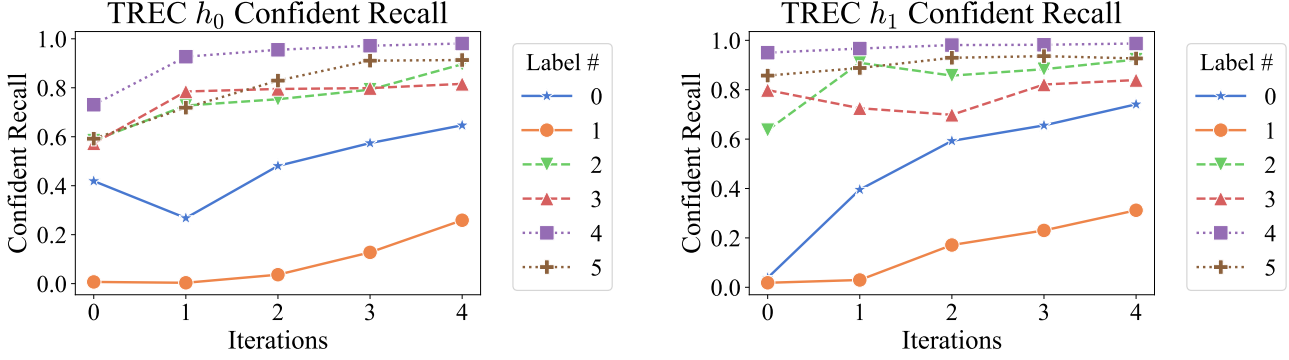


Figure 7. Recall of the confident pseudolabel set L_0^t (left, extracted from h_0 using Algorithm 2) and L_1^t (right, extracted from h_1 using Algorithm 3) for each label versus co-training iteration t .

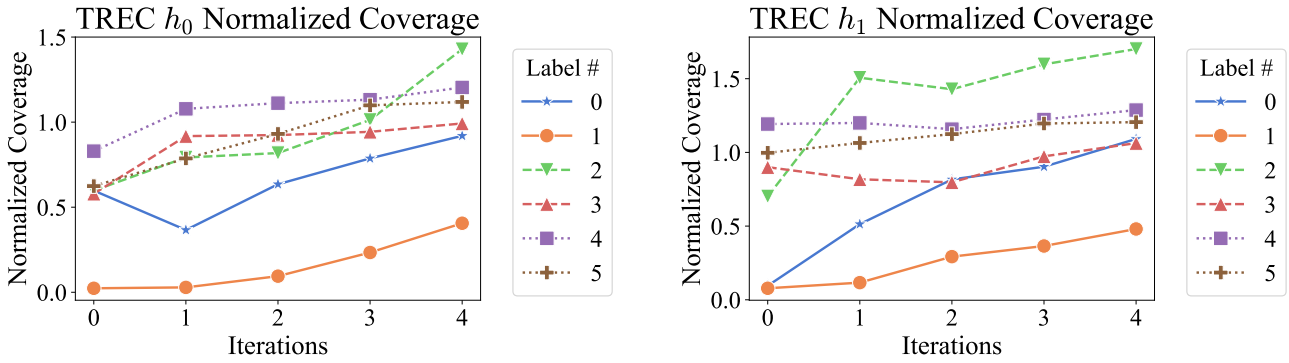


Figure 8. Normalized coverage of the confident pseudolabel set L_0^t (extracted from h_0 using Algorithm 2) for each label versus co-training iteration t . Normalized coverage for label j is computed as $|\{(x, \hat{y}) \in L_i^t : \hat{y} = j\}| / |\{x : y(x) = j\}|$ (the number of examples with confident pseudolabel j divided by the number of examples with true label j). This metric decouples the coverage from the precision. The increasing slope of label 1 (left) indicates that h_0 adds more confident data for label 1 in the later iterations. Combining this with the label 1 precision versus iteration curve in Figure 9 (left) indicates that the model adds more confident data for label 1 as it gets more precise, which is the desired behavior. On the other hand, for other labels (e.g. label 4) the rate of confident data addition and the precision stay relatively constant.

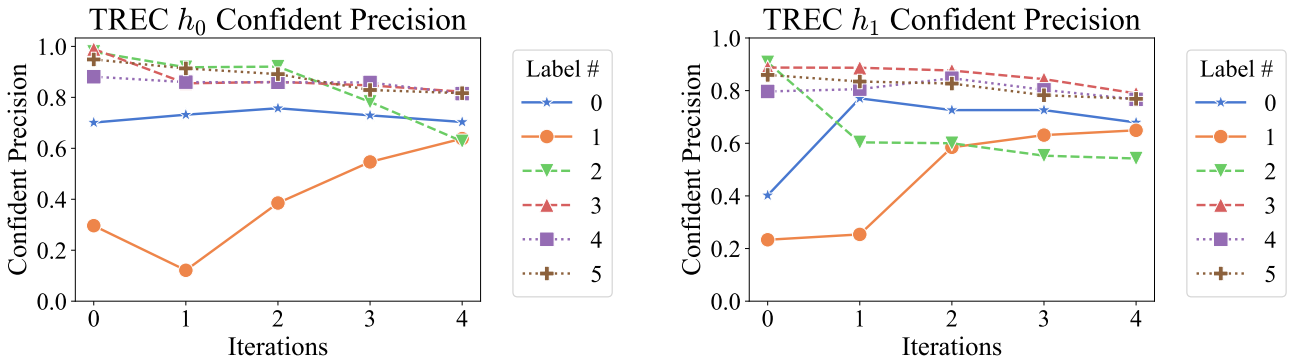


Figure 9. Precision per label vs co-training iteration, h_0 (left—identical to right display of Figure 3) and h_1 (right). Together with Figure 8, this indicates the two regimes of co-training. For labels 0 and 2-5, the precision decreases or remains the same while the coverage increases roughly linearly. This is the “coverage-expansion” regime, where the initial confident pseudolabels are high-precision and the model learns to imperfectly extend that initial signal to the uncovered data with some losses in precision. This regime is present in classical co-training results (Balcan et al., 2005). On the other hand, for label 1, both the coverage and the precision increase with the iteration t . This is the *pseudolabel-correction* regime, because the models are able to learn to be more accurate than the pseudolabels used to train them (compare h_0 precision for label 1 to h_1 precision for label 1 in the same iteration—the h_1 model is trained on the labels from h_0 , but is able to select confident data with better precision than those labels).

D. Prompts

Section D.1 discusses the impact of prepending the [PAD] token *before* obtaining initial labels from T0-3B. The following two subsections list the prompts used for our experiments, which are largely taken from Sanh et al. (2022).

D.1. Prepending the [PAD] token

As described in Section 3, for the first step of co-training T0-3B, we initialize the soft prompt with the [PAD] token, then generate the initial pseudolabels with these tokens prepended to the input. We made this choice for conceptual consistency, so that the model trained in the first round of co-training is the same model that was used to generate the pseudolabels. However, surprisingly, we found that this addition of 20 [PAD] tokens before the input actually sometimes improved the performance of zero-shot prompting. Table 3 compares the performance of T0-3B zero-shot with and without the [PAD] tokens prepended.

Dataset	T0-3B ([PAD])	T0-3B (no [PAD])
RTE	68.9	62.5
CB	58.9	55.4
BoolQ	56.4	59.0

Table 3. Comparison between zero-shot prompt performance with T0-3B with and without 20 [PAD] tokens prepended to the input. We report numbers for the prompts from Section D.3, but the performance differences are similar for other prompts as well.

D.2. Partial Access Setting Prompts

RTE

```
{example_premise}
Question: {example_hypothesis} True, False,
        or Unknown?
answer: {example_answer}
{premise}
Question: {hypothesis} True, False, or
        Unknown?
answer:
```

CB

```
hypothesis: {example_hypothesis}
Does the premise imply the hypothesis? Yes,
        No, or Neither?
answer: {example_answer}
premise: {premise}
hypothesis: {hypothesis}
Does the premise imply the hypothesis?
Yes, No, or Neither?
answer:
```

TREC

```
Classify the questions based on whether
        their answer type is Unknown, Number,
        Location, Person, Description, Entity,
        or Abbreviation.
Question: {example_question}
Answer Type: {example_type}
Question: {question}
Answer Type:
```

D.3. Full Access Setting Prompts

RTE

```
{premise}
Question: {hypothesis} True, False, or
        Unknown?
```

CB

```
{premise}
Question: {hypothesis} True, False, or
        Neither?
```

BoolQ

```
Text: {passage}
Answer the following yes/no question: {
        question}? Yes or no?
```

Algorithm 2 GetConfDataMC

```

input  $\{x_n\}_{n=1}^U$  unlabeled examples
input model  $h_i$ , view  $\phi_i$ 
input coverage fraction  $\tilde{\beta}$ 
input minimum class percentage  $\gamma$ 
    // compute pseudolabel and score for each example
    for  $n$  in  $\{1, \dots, U\}$  do
         $o_n = h_i(\phi_i(x_n))$  (note  $o_n \in \mathbb{R}^{numlabels}$ )
         $\hat{y}_n \leftarrow \operatorname{argmax}_l o_{nl}$ 
         $s_n \leftarrow \max_l o_{nl}$ 
    end for
     $L \leftarrow \emptyset$ 

    // first, select top  $\gamma\%$  for each class by score
     $ms \leftarrow \lfloor \gamma \tilde{\beta} U \rfloor$  // min num points to select
    for  $l$  in  $\{1, \dots, numlabels\}$  do
         $\mathcal{I}_l = \{(x_n, \hat{y}_n, s_n) : \hat{y}_n = l\}$ 
        // sort  $\mathcal{I}_l$  by score  $s_n$  (ascending)
         $S_l \leftarrow \operatorname{Sort}(\mathcal{I}_l, \text{key}=\text{lambda } q: q[2])$ 
        // add top  $\alpha\%$  to  $L$  (read off end of  $S_l$ )
         $L \leftarrow L \cup S_l[-ms:]$ 
    end for

    // now select the rest of the points
     $rs \leftarrow \lfloor \tilde{\beta} U \rfloor - |L|$  // num remaining points to select
     $\mathcal{I} = \{(x_n, \hat{y}_n, s_n)\}_{n=1}^U$ 
     $\mathcal{I} \leftarrow \mathcal{I} \setminus L$  // don't select twice
     $S \leftarrow \operatorname{Sort}(\mathcal{I}, \text{key}=\text{lambda } q: q[2])$ 
     $L \leftarrow L \cup S[-rs:]$ 

    // chop off score and return point + pseudolabel
     $L \leftarrow \{(x_n, \hat{y}_n) | \exists s_n : (x_n, \hat{y}_n, s_n) \in L\}$ 
return  $L$ 

```

Algorithm 3 GetConfDataCS

```

input  $\{x_n\}_{n=1}^U$  unlabeled examples
input model  $h_i$ , view  $\phi_i$ 
input coverage fraction  $\tilde{\beta}$ 
    // compute pseudolabel and repr. for each example
    for  $n$  in  $\{1, \dots, U\}$  do
         $o_n = h_i(\phi_i(x_n))$  (note  $o_n \in \mathbb{R}^{numlabels}$ )
         $\hat{y}_n \leftarrow \operatorname{argmax}_l o_{nl}$ 
    end for
     $L \leftarrow \emptyset$ 
    for  $y$  in  $\{1, \dots, numlabels\}$  do
         $\hat{P}_y = |\{n : \hat{y}_n = y\}|/U$ 
    end for

    // compute 20 nearest neighbors for each ex.
    for  $u$  in  $\{1, \dots, U\}$  do
         $N(u) = \operatorname{NN}_{20}(\phi_i(x_u), \{\phi_i(x_v)\}_{v=1}^U)$ 
        for  $v$  in  $N(u)$  do
             $w_{uv} = 1/(1 + \|\phi(x_u) - \phi(x_v)\|_2)$ 
             $I_{uv} = \mathbb{I}[\hat{y}_u \neq \hat{y}_v]$ 
        end for
        // now compute cut statistic
         $J_u = \sum_{v \in N(u)} w_{uv} I_{uv}$ 
         $\mu_u = (1 - \hat{P}_{\hat{y}_u}) \sum_{v \in N(u)} w_{uv}$ 
         $\sigma^2 = \hat{P}_{\hat{y}_u} (1 - \hat{P}_{\hat{y}_u}) \sum_{v \in N(u)} w_{uv}^2$ 
         $s_u = \frac{J_u - \mu_u}{\sigma}$ 
    end for
    // now sort by statistic and return top data
     $\mathcal{I} \leftarrow \{(x_n, \hat{y}_n, s_n)\}_{n=1}^U$ 
     $S \leftarrow \operatorname{Sort}(\mathcal{I}, \text{key}=\text{lambda } q: q[2])$ 
     $ns = \lfloor \tilde{\beta} U \rfloor$ 
     $L \leftarrow S[:ns]$ 
     $L \leftarrow \{(x_n, \hat{y}_n) | \exists s_n : (x_n, \hat{y}_n, s_n) \in L\}$ 
return  $L$ 

```

Algorithm 4 Co-training algorithm (detailed, GPT-3)

```

input  $\{(x_j, y_j)\}_{j=1}^k$  initial labeled examples
input  $\mathcal{U} = \{x_n\}_{n=1}^U$  unlabeled examples
input initial coverage  $\beta$ , coverage increase  $\beta'$ 
input minimum percentage per class  $\gamma$ 
    // build view 0 for unlabeled examples
    for  $n$  in  $\{1, \dots, U\}$  do
        for  $j$  in  $\{1, \dots, k\}$  do
             $\phi_0^{(j)}(x_n) \leftarrow \text{GPT3}(x_n, (x_j, y_j))$ 
        end for
         $\phi_0(x_n) \leftarrow (\phi_0^{(1)}; \dots; \phi_0^{(k)})$ 
    end for

    // build view 1 for unlabeled examples
    for  $n$  in  $\{1, \dots, U\}$  do
        // extract pre-trained DeBERTa representation for  $x_n$ 
         $\phi_1(x_n) \leftarrow \text{DeBERTa}(x_n)$ 
    end for

    // initialize  $h_0$  according to (1)
    for  $j$  in  $\{1, \dots, k\}$  do
        // get GPT-3 outputs on content-free input
         $\phi_0^{(j)}(x_{cf}) \leftarrow \text{GPT3}(x_{cf}, (x_j, y_j))$ 
         $W^{(j)} \leftarrow \text{Diag}\left(\frac{1}{\phi_0^{(j)}(x_{cf})}\right)$ 
    end for
     $W \leftarrow \{W^{(j)}\}_{j=1}^k$ 
     $\alpha \leftarrow \mathbf{1}$ 
     $h_0 \leftarrow h_0(\cdot; W, \alpha)$ 

    // co-training loop
    for  $t$  in  $\{0, \dots, T-1\}$  do
         $\tilde{\beta} \leftarrow \beta + t\beta'$ 
         $L_0^t \leftarrow \text{GetConfDataMC}(\mathcal{U}, h_0, \phi_0, \tilde{\beta}, \gamma)$ 
         $h_1 \leftarrow \text{Train}(\phi_1, L_0^t)$ 
         $L_1^t \leftarrow \text{GetConfDataCS}(\mathcal{U}, h_1, \phi_1, \tilde{\beta})$ 
         $h_0 \leftarrow \text{Train}(\phi_0, L_1^t)$ 
    end for
return  $(h_0, h_1)$ 

```

Algorithm 5 Co-training algorithm (detailed, T0)

```

input  $\mathcal{U} = \{x_n\}_{n=1}^U$  unlabeled examples
input initial coverage  $\beta$ , coverage increase  $\beta'$ 
    // format the input with a hard prompt template  $P$ ,
    // then get T0 embedding to build view  $\phi_0$ 
    for  $n$  in  $\{1, \dots, U\}$  do
         $\tilde{x}_n \leftarrow P(x_n)$ 
         $\phi_0(x_n) \leftarrow \text{T0Emb}(\tilde{x}_n)$ 
    end for

    // build view 1 for unlabeled examples
    for  $n$  in  $\{1, \dots, U\}$  do
        // extract pre-trained DeBERTa representation for  $x_n$ 
         $\phi_1(x_n) \leftarrow \text{DeBERTa}(x_n)$ 
    end for

    // initialize soft prompt with repeated pad token emb
     $p \leftarrow \text{T0Emb}([\text{PAD}])$ 
     $h_0(\cdot) \leftarrow \text{T0}((p; p; \dots; p); \cdot)$ 

    // co-training loop
    for  $t$  in  $\{0, \dots, T-1\}$  do
         $\tilde{\beta} \leftarrow \beta + t\beta'$ 
         $L_0^t \leftarrow \text{GetConfDataCS}(\mathcal{U}, h_0, \phi_0, \tilde{\beta})$ 
         $h_1 \leftarrow \text{Train}(\phi_1, L_0^t)$ 
         $L_1^t \leftarrow \text{GetConfDataCS}(\mathcal{U}, h_1, \phi_1, \tilde{\beta})$ 
         $h_0 \leftarrow \text{Train}(\phi_0, L_1^t)$ 
    end for
return  $(h_0, h_1)$ 

```
