A TRANSFER LEARNING FRAMEWORK FOR WEAK TO STRONG GENERALIZATION

Seamus Somerstep^{*} Felipe Maia Polo^{*} Moulinath Banerjee^{*} Ya'acov Ritov^{*} Mikhail Yurochkin^{**} Yuekai Sun^{*} ^{*} University of Michigan

* IBM Research * MIT-IBM Watson AI Lab

Abstract

Modern large language model (LLM) alignment techniques rely on human feedback, but it is unclear whether these techniques fundamentally limit the capabilities of aligned LLMs. In particular, it is unknown if it is possible to align (stronger) LLMs with superhuman capabilities with (weaker) human feedback *without degrading their capabilities*. This is an instance of the weak-to-strong generalization problem: using feedback from a weaker (less capable) model to train a stronger (more capable) model. We prove that weak-to-strong generalization is possible by eliciting latent knowledge from pre-trained LLMs. In particular, we cast the weakto-strong generalization problem as a transfer learning problem in which we wish to transfer a latent concept prior from a weak model to a strong pre-trained model. We prove that a naive fine-tuning approach suffers from fundamental limitations, but an alternative refinement-based approach suggested by the problem structure provably overcomes the limitations of fine-tuning. Finally, we demonstrate the practical applicability of the refinement approach in multiple LLM alignment tasks.

1 INTRODUCTION

Modern AI alignment methods are based on human feedback, but such methods may limit the abilities of AI models to those of human experts. When the capabilities of AI systems exceed those of humans (Bengio, 2023), human experts may not be able to comprehend–much less provide feedback on–the outputs of AI models. For example, future AI models may be able to develop entire software stacks in multiple programming languages that no (human) software engineer can review in their entirety. This leads to the superalignment problem (Leike & Sutskever, 2023): aligning superhuman AI when human experts can only provide (relatively) weak guidance.

Following Burns et al. (2023a), we study superalignment through the analogy of training more capable models (*i.e.*, GPT-4o-mini) on outputs from weaker models (*i.e.*, Llama-7B). This problem setting, using a smaller weaker model (instead of humans) to supervise the alignment of a larger stronger model, is known as *weak to strong generalization* (Burns et al., 2023a). Our main contributions are:

- 1. We formulate the weak-to-strong generalization problem as a transfer learning problem in which we wish to transfer a prior over a latent concept from a weaker to a stronger model.
- 2. Within our framework, we show that naively fine-tuning the strong model with the weak labels leads to an estimate for the target function with poor expected risk. Empirically, we demonstrate that the accuracy of the fine-tuned strong model is limited by the accuracy of the weak model because the strong model will learn to emulate the mistakes of the weak model.
- 3. Motivated by these negative results, we develop a refinement-based approach that elicits latent knowledge from the strong model. Within our framework, we model this as an implicit Bayesian inference step and prove that it overcomes the limitations of fine-tuning on the weak labels. We also demonstrate the practical applicability of this approach by helping GPT-3.5-Turbo (Brown et al., 2020) and GPT-40-mini (OpenAI, 2024) learn a new persona, improve mathematical reasoning, and learn an explanation technique with weak supervision provided by a variety of weak models.

^{*}smrstep@umich.edu

1.1 CONCURRENT WORK

We reserve the main paper for a discussion of closely related ideas in the weak-to-strong generalization space. Please see Appendix G for prior related work.

Simultaneously to us, the authors of Lang et al. (2024), Charikar et al. (2024) and Wu & Sahai (2024) have developed their own theories of weak to strong generalization. In Charikar et al. (2024), a representation based model is proposed. Under the assumption that models are selected over a convex set, they quantify the gain of the weak-label trained strong model over the weak model. The authors of Lang et al. (2024) show that weak-to-strong generalization will arise if the strong model has good properties over neighborhoods of the data space. The work of Wu & Sahai (2024) demonstrates that weak to strong generalization can occur in spiked covariance models with growing dimensions and label space cardinality. Our theory is generally distinct; each of the mentioned works is concerned with showing that simply fine-tuning on weak labels is able to provide substantial gains on the target task (as compared to the weak model). The primary message of our work is different. In our formulation, the benefit of weak label training is quite limited, and thus we provide alternative procedures to work around this issue.

Empirically, the closest work to ours is Yang et al. (2024b). Concurrently they develop a methodology similar to our refinement method. Our work provides statistical intuition for the success of their methods, and as part of our experimental contribution, we show that our refinement method achieves weak-to-strong generalization on their weakly labeled data sets. Other emprical investigations of weak to strong generalization include Zhang et al. (2024) which investigates the role of the temperature parameter in weak to strong generalization, Yang et al. (2024a) which studies deception in weak to strong generalization, and Fan et al. (2024) which proposes a dynamic logit fusion approach for weak to strong generalization.

2 THE TRANSFER LEARNING MODEL

Our transfer learning framework is built on the varous mixture and/or latent concept models for large language models (Xie et al., 2021; Wang et al., 2024; Pathak et al., 2024). Generally, these models hypothesize that an LLM fed a prompt of the form $prompt = ((x_1, y_1), (x_2, y_2), \ldots, (x_{n_{\text{ICL}}}, y_{n_{\text{ICL}}}), x)$, has distribution

$$P(y|\text{prompt}) = \sum_{k} p(y|x,k) p(k|(x_1,y_1),(x_2,y_2),\dots,(x_{n_{\text{ICL}}},y_{n_{\text{ICL}}}),x).$$
(2.1)

2.1 SET UP

Motivated by the weak-to-strong alignment problem (Burns et al., 2023a), we consider the following transfer learning problem: There are source and target distributions P and Q which are both joint distributions on the tuple of random variables $(X, Y, Y') \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}$, *i.e.* we have

$$(X, Y, Y') \sim P$$
 or Q .

In the transfer learning problem, each of the random variables in this tuple represents a different aspect of the weak-to-strong generalization problem. The variable X represents (tokenized) queries to an LLM, and the variables Y and Y' represent the outputs of the strong and weak models, respectively. Consequently, $P_{Y|X}$ will represent a strong unaligned LLM fed X while $Q_{Y'|X}$ will represent a weaker but aligned LLM fed an input X. The ultimate goal of the learner is to produce an aligned version of the strong model (in our set-up, this is represented as $Q_{Y|X}$).

Source distribution: In practical weak to strong generalization settings, the learner often does not observe direct samples from P but merely has access to $P_{Y|X}$ (this represents a practitioner who is given a model that must be aligned). Furthermore, the learner does not observe weak labels in the source distribution, making the specification of $P_{Y'|X}$ irrelevant. We turn our attention to the pairs (X, Y) and assume that we may write

$$P_X \stackrel{d}{=} \mathrm{Unif}([-1,1]^d); \quad P_{Y|X} \stackrel{d}{=} \sum_{k=1}^K \alpha_k \varphi(y; \beta_k^\top X, \sigma^2).$$

Here φ denotes the normal density, σ^2 the error rate on the labels, and $\sum_k \alpha_k = 1$ with $\alpha_k \ge 0$ for all k. The choice to use a linear model is influenced by the results in Pathak et al. (2024) where

transformer architectures that emulate mixtures-of-regressions are studied. The reader may also interpret this set up as the one studied in Wang et al. (2024) with a linear specification on their map between X and Y. The connection between the framework and these references is discussed further in Appendix A. We emphasize that the concepts k and model components β_k are latent; the learner can view Y for a given X but does not know which internal component generated Y.

Target distribution: In weak-to-strong generalization, the learner does not observe gold-standard target data $((X, Y) \sim Q_{X,Y})$. Instead, they are given covariates X and (possibly biased or noisy) labels Y' that are produced by a weak(er) model fed X. Following Wang et al. (2024) we assume that the target data is drawn from the source distribution with a prior shift towards one desirable concept. In other words, we may write the conditional distributions of Y and Y' as the following:

$$Q_{Y|X} \stackrel{d}{=} \varphi(y; \beta_{k^*}^\top X, \sigma^2); \quad Q_X \stackrel{d}{=} \text{Unif}([-1, 1]^d)$$
$$Q_{Y'|X} \stackrel{d}{=} \begin{cases} \varphi(\beta_{k^*}^{w^\top} X; \sigma^2) & \text{biased weak model,} \\ \varphi(\beta_{k^*}^\top X; \sigma^2 + {\sigma'}^2) & \text{noisy weak model.} \end{cases}$$

In other words, the weak model provides target supervision in the sense that it is sampling from the correct concept, but the conditional density is corrupted, representing the reduced capabilities of smaller language models. In this view, the prior over the latent concept k represents alignment, while the corrupted density represents the weakened capabilities of the smaller aligned model. Throughout, we will need minor assumptions on the vectors β_k and $\beta_{k^*}^w$.

Assumption 2.1. The collection of vectors $\{\beta_k\}_{k=1}^K$ is orthonormal. Additionally, the vector $\beta_{k^*}^w$ satisfies $\begin{cases} \beta_{k^*}^{w^\top} \beta_k \in [0,1], & \text{if } k = k^* \\ \beta_{k^*}^{w^\top} \beta_k = 0, & \text{if } k \neq k^* \end{cases}$

We have opted to consider two versions of weakness, the noisy case is simply if the weak labels are the strong labels corrupted by iid noise. In the other case, the weak labels are provided by a model with some misspecified parameter $\beta_{k^*}^w$. The quantity $\beta_{k^*}^T \beta_{k^*}^w$ measures how closely the weak aligned model approximates the target model.

The ultimate goal of the learner is to take data, $D' = \{x_i, y'_i\}_{i=1}^{n_{Q'}}$, (sampling) access to $P_{Y|X}$ and produce an estimator $\hat{\beta}$ that predicts Y from X over Q. We are interested in the excess risk of any produced estimate $\hat{\beta}$ which is given by $\mathcal{R}(\hat{\beta}) = \mathbb{E}_{Q_X}[\beta_{k^*}^\top X - \hat{\beta}^\top X]^2$.

The following is an example of our framework in practice, here we wish to teach a strong model, *i.e.* one that produces factually correct (uncorrupted) responses, to talk like a pirate. The source concept prior is the distribution of personas of a stereotypical LLM, while the target concept is the pirate persona. In weak to strong generalization, a weak model produces outputs that are corrupted, but from the target concept. Here Falcon7B is explicitly instructed to respond to questions as a pirate, one can see that the resulting label is a response in a pirate style that is factually incorrect (Paul Newman played Billy the Kid in the film The Left Handed Gun).

Example 2.2 (Persona Learning).

 $\vec{\alpha}^p$: The source prior is the standard personas of an AI.

 k^* : The target concept is characterized by a pirate persona.

X:"Who played Billy the Kid in the Left Handed Gun?"

 $\mathit{Falcon7B}(Q_{Y'|X}){:}"Ahoy, me hearties! Billy the Kid was played by the legendary actor, John Wayne."$

3 DIFFICULTY AND FEASIBILITY OF WEAK TO STRONG GENERALIZATION

In our transfer learning setup, we made a particular and non-standard assumption on the relationship between the source conditional $P_{Y|X}$ and the target $Q_{Y|X}$ conditional distributions. In particular,

 $\mathbb{E}_P[Y|X]$ is specified by a mixture of functions $\beta_1 \dots \beta_K$ and $\mathbb{E}_Q[Y|X]$ remains in the convex hull of the source mixture. In the next two sections, we demonstrate two ideas: First, because of the lack of any strong target supervision, the problem is intractable without some structure on the relationship between the source and the target. Second, we demonstrate that our convex hull assumption is sufficient for the learner to improve the weak supervision, allowing for weak to strong generalization.

3.1 LIMITATIONS OF TRAINING ON THE WEAK LABELS

In this (sub)-section we consider (one of) the standard methods for achieving weak to strong generalization proposed by Burns et al. (2023a), that is, to train the source model on the weak labels, with some shrinkage towards the source model. In our regression setting, we study a family of "shrinkage to source" estimators which have the following form.

Definition 3.1. We define the weakly trained estimator β_n as the estimator that satisfies the following

$$\hat{\beta}_{\eta} \triangleq \arg\min_{\beta \in \mathbb{R}^d} \frac{1}{n_{Q'}} \sum_{i=1}^{n_{Q'}} [y'_i - \beta^T x_i]^2 - \frac{\eta}{2} ||\beta - (\sum_k \alpha_k \beta_k)||^2$$

The first term in definition 3.1 rewards the model for generating responses that approximate the weak labels, while the second term represents that fine-tuning is often performed with some form of regularization towards the source. In the weak-to-strong generalization problem, this term represents the fact that only a portion of the model weights is altered during fine-tuning and only for a limited number of epochs. In true superalignment, a KL-divergence-based regularization term is often explicitly encoded in the training objective, for example, if RLHF is used for the alignment procedure (Ouyang et al., 2022). Perhaps unsurprisingly, the expected MSE of $\hat{\beta}_{\eta}$ is generally poor.

Proposition 3.2. Consider the case of biased weak supervision; if ϵ_P^2 and $\epsilon_{Q'}^2$ denote the squared bias of the source and weak models, then the following lower bound on the MSE of estimators produced by naive fine-tuning holds:

$$\mathbb{E}_{Q'}\mathcal{R}(\hat{\beta}_{\eta}) \ge \left[\frac{1}{1+\eta}\right]^2 \epsilon_{Q'}^2 + \left[\frac{\eta}{1+\eta}\right]^2 \epsilon_P^2 + \frac{\eta}{(1+\eta)^2} (1-\alpha_{k^*})(1-\beta_{k^*}^T \beta_{k^*}^w)$$

The intuition behind this result is that both the source of weak labels and the source model are biased estimators of the target model; this leads to an impossibility result where we can't simply combine the two to create a consistent estimator. Propositon 3.2 is somewhat in contrast with prior work on learning under weak supervision, self-training, and label propagation (Vishwakarma et al., 2022; Shin et al., 2023; Ratner et al., 2020). In many of these scenarios, some form of training on weak labels can be beneficial to the learner. Here, the learner only has access to weak labels from one weak teacher, and despite the latent cluster assumption, the outputs are generally continuous and thus do not satisfy the well-separability assumption often seen in self training (Grandvalet & Bengio, 2004). Finally, while we have only considered the case of biased weak supervision in Proposition 3.2, in the noisy case the estimator $\hat{\beta}_{\eta}$ is unbiased if $\eta = 0$, however, it does have a poor asymptotic relative efficiency compared to an estimator fit on gold standard labels. As an example, consider the following query of GPT3.5 after it is fine-tuned on pirate-themed responses provided by Falcon 7B. Note that in particular, the weakly trained version of GPT 3.5 struggles with a simple question that presents no issue for the unaltered version.

Example 3.3 (Persona Learning Test). Consider the following example of sampling a test response from GPT3.5, after it is fine-tuned on responses from Falcon7B.

X:"Come up with words that rhyme with the given word: Instruct " $GPT3.5(P_{Y|X})$:"Here are some words that rhyme with "instruct": Conduct, Construct, Destruct, ..." Weakly trained GPT3.5 ($\hat{\beta}_{\eta}$):"Ahoy, me hearties! I'll be instructin' ye to come up with words that rhyme with the given word. *winks*"

3.2 GEOMETRIC INTUITION FOR THE FEASIBILITY OF WEAK TO STRONG GENERALIZATION

We have seen that the difficulty of weak to strong generalization arises from the poor quality of the weak target supervision. Fortunately, our transfer learning structure suggests a solution to the problem: we must utilize the fact that the target function $\mathbb{E}_Q[Y|X]$ is contained in the convex hull of the source model. In practice, this means that the source model has the latent knowledge to complete the target task; it just needs this ability unlocked by utilizing the weak supervision.

Recall that we have access to a weakly labeled data set $D' = (\mathbf{X}, \mathbf{Y}')$. To see the intuition behind the proposed methods in the paper, imagine that the learner has actual access to the collection of prediction vectors (over \mathbf{X} in the weakly labeled data) from each component of the source model. We can write this collection as $F \triangleq \mathbb{E}_P[\mathbf{Y}|\mathbf{X}, k]_{k=1}^K = [\beta_1^T \mathbf{X} | \cdots | \beta_K^T \mathbf{X}] \in \mathbf{R}^{n_{Q'} \times K}$. The learner may opt to produce new labels \hat{y} by solving

$$\begin{cases} \hat{y} \leftarrow F\hat{w} \\ \hat{w} \leftarrow \arg\min_{w \in \Delta^{K-1}} \frac{1}{2} \|y' - Fw\|_2^2 \end{cases} \equiv \hat{y} \leftarrow \arg\min_{g \in \operatorname{cvx}(F)} \frac{1}{2} \|y' - g\|_2^2, \qquad (3.1)$$

where $\operatorname{cvx}(F) \subset \mathbb{R}^{n_{Q'}}$ is the convex hull of $\beta_1^T \mathbf{X} \dots \beta_K^T \mathbf{X}$. We now show that our convex hull assumption on the source and target distribution is sufficient for the possibility of weak-to-strong generalization.

Proposition 3.4. Define $\epsilon' \triangleq \mathbf{Y}' - \mathbb{E}_Q[\mathbf{Y}|\mathbf{X}] \in \mathbb{R}^{n_{Q'}}$. If $\mathbb{E}_Q[\mathbf{Y}|\mathbf{X}] \in \operatorname{cvx}(F)$, then \hat{y} in (3.1) satisfies

$$\mathbb{E}\left[\frac{1}{n_{Q'}}\|\hat{y} - \mathbb{E}_Q[\mathbf{Y}|\mathbf{X}]\|_2^2\right] \le \frac{1}{n_{Q'}}\mathbb{E}\left[\sup_{\theta \in T_{\mathrm{cvx}(F)}(\mathbb{E}_Q[\mathbf{Y}|\mathbf{X}]) \cap \mathbf{S}^{n-1}}(\epsilon'^\top \theta)^2\right] \ll \frac{1}{n_{Q'}}||\epsilon'||^2,$$

where $T_{\text{cvx}(F)}(\mathbb{E}_Q[\mathbf{Y}|\mathbf{X}])$ is the tangent cone of cvx(F) at $\mathbb{E}_Q[\mathbf{Y}|\mathbf{X}]$.

We wish to emphasize multiple aspects of this observation. First, this is merely an analogy for eliciting knowledge from the strong model to improve the weak supervision; in practice, the learner does not set up and solve (3.1) (recall that the learner does not even have access to the actual mixture components). Instead, the learner feeds the weakly labeled data to the strong model for refinement. We formalize this in Section 4. Second, Proposition 3.4 is a statement on the quality of the supervision *on the training data*, the learner will still need to fit a model to x, \hat{y} . One may also note that the right side of the first inequality in Proposition 3.4 is a geometric complexity measure called the statistical dimension (Amelunxen et al., 2014), and it is a key quantity in the study of statistical efficiency in high dimensions.

4 WEAK TO STRONG GENERALIZATION WITH OUTPUT REFINEMENT

In Section 3.2 we saw that as long as the convex hull assumption holds and the learner has access to both the source weights α_k and the components β_k , weak supervision over the target can be improved by leveraging the source information to produce a set of refined labels $\hat{\mathbf{Y}}$ for \mathbf{X} . At first blush, this may seem unhelpful for aligning a complex LLM, since directly accessing the components or weights is not possible. Fortunately, two basic ideas will allow us to execute 3.1 in practice: First, if we draw the refined label from the unaligned strong model, the label is guaranteed to be in the same convex hull as the hypothetical strong target labels. Second, the learner is able to indirectly manipulate the weights through a combination of in-context learning examples and an optional system prompt.

4.1 **REFINEMENT WITH IN-CONTEXT-LEARNING**

To steer the weights of the unaligned LLM, we follow the philosophy of Wang et al. (2024) and propose that the learner utilize the implicit Bayesian inference capabilities of an LLM.

Formally, consider a prompt x for which we wish to obtain better supervision. To do so, we select ICL examples $S_{n_{\rm ICL}} = \{(x_j, y'_j)\}_{j=1}^{n_{\rm ICL}}$ from the weakly labeled training data set, form a concatenated prompt $[S_{n_{\rm ICL}} \circ x]$, and re-sample a new label from the source model fed the concatenated prompt. In practice, we have a finite weakly labeled data set $D' : \{(x_i, y'_i)\}_{i=1}^{n_{Q'}}$, from which we will attain both the training questions and the ICL examples. Algorithm 1 summarizes this procedure.

Algorithm 1 ICL Refinement

Require: Weakly labelled data $D' : \{(x_i, y'_i)\}_{i=1}^{n_{Q'}}$, source LLM. 1: Select ICL examples $S_{n_{\text{ICL}}} \leftarrow \{(x_j, y'_j)\}_{j=1}^{n_{\text{ICL}}}$ 2: Set aside remaining training data $\{x_i, y'_i\}_{i=1}^{n_{\text{Re}}} \leftarrow D' \setminus S_{n_{\text{ICL}}}$

3: for $i = 1, 2, ..., n_{re}$ do

4: Construct concatenated prompt $[S_{n_{\text{ICL}}} \circ x_i]$,

5: Draw refined label \hat{Y} from the source LLM fed the concatenated prompt.

6: end for

7: return $\hat{\mathcal{D}} = \{x_i, \hat{y}_i\}_{i=1}^{n_{\text{re}}}$

Recall the setting of Example 2.2, where we wish to train an advanced model to use a new persona. The weak responses provided are generally in the correct persona but contain factual errors. To correct this issue, we utilize the capable model to infer the correct concept from the weak labels (using in-context learning) to provide better labels. Here is an example of a resampled label using the ICL method. The style has been inferred from the weak labels, but since we are sampling from the stronger model, the labels are now factually correct. The reader should compare this with the quality of the label in Example 2.2.

Example 4.1 (Persona Learning Label Re-sample). The following is an example of a response resampled from GPT3.5 (GPT 3.5 is also fed other weakly labeled instances as ICL examples).

> X: "Who played Billy the Kid in the Left Handed Gun?" GPT3.5[ICL + X] $(P_{\hat{Y}|X})$:"Ahoy, me hearties! In the film "The Left Handed Gun," Billy the Kid was played by none other than Paul Newman. *winks*"

4.1.1 THEORETICAL ANALYSIS OF ICL REFINEMENT

Suppose we have our fixed design matrix X over which we are interested in attaining improved supervision for, and to do so we will select ICL examples $(S'_{n_{ICL}})$, and re-sample a label from the source model fed the ICL examples plus the training query X. Denoting the refined labels as \hat{Y} , we will ultimately fit the estimator

$$\hat{\beta}_{\text{re}} = \arg\min_{\beta \in \mathbb{R}^d} \sum_{i=1}^{n_{\text{re}}} ||\hat{y}_i - \beta^T x_i||^2.$$

The key to studying the quality of $\hat{\beta}_{re}$ is controlling the quality of the refined labels \hat{Y} . Ultimately, we are interested in quantifying the penalty the learner suffers from feeding the weakly labelled ICL examples to the source model (as compared to using hypothetical gold standard labels). To do this, we must specify the form that $P(k|S_{n_{ICL}})$ takes. There are two cases we consider, the first is where the model treats the examples as iid.

Assumption 4.2 (iid ICL examples). The following distributional assumption holds on the refined labels \hat{Y} :

$$P_{\hat{Y}|X} \stackrel{d}{=} \sum_{k} \frac{\alpha_{k} \prod_{j=1}^{n_{lCL}} P_{y'_{j}|x_{j},k}}{\sum_{k'} \alpha_{k'} \prod_{j=1}^{n_{lCL}} P_{y'_{j}|x_{j},k'}} \varphi(\beta_{k}^{T}X;\sigma^{2})$$

$$\stackrel{d}{=} \sum_{k} \frac{\alpha_{k} e^{-\sum_{j=1}^{n_{lCL}} (y'_{j} - x_{j}^{T}\beta_{k})^{2}}}{\sum_{k'} \alpha_{k'} e^{-\sum_{j=1}^{n_{lCL}} (y'_{j} - x_{j}^{T}\beta_{k'})^{2}}} \varphi(\beta_{k}^{T}X;\sigma^{2}).$$

The assumption that the model is treating the ICL examples as iid is an admittedly strong one. Despite this, there are two settings in the literature where this holds (further discussion on this is supplied in Appendix A). First, is in the model proposed by Wang et al. (2024) (here we make further specification on the relationship between X and Y). Second, is if the architecture proposed in Pathak

et al. (2024) produces the refined labels. Beyond this iid assumption, we provide an analysis in the setting of Xie et al. (2021) in Appendix B.

We pause to consider what the above equation represents: we are assuming the source model selects k based on its own implicit beliefs on the relationship between y' and x, then produces a new label. The learner is not actually calculating $P(k|S_{n_{ICL}})$, rather the base LLM is using its in-context-learning capabilities to infer the target concept prior from the weakly labeled data. This raises a potentially sticky issue: the source model is unaware that the ICL examples are drawn from some potentially misspecified model, and as such it simply evaluates the likelihood of the latent concept k, assuming that they are drawn from the correctly specified regression function. The key question is as follows: Can the correct concept be inferred from the ICL examples despite the fact that they are not the gold standard? The following theorem addresses this issue in our setting.

Theorem 4.3. Suppose that both the orthonormality and inner product assumptions on $\{\beta_k\}$, $\beta_{k^*}^w$ of the transfer learning framework hold. Additionally, assume that \hat{Y} is drawn from the distribution in Assumption 4.2. Then in both the case of biased or weak supervision, the following finite sample bound on the excess risk of $\hat{\beta}_{re}$ holds:

$$\mathbb{E}_{Q',\hat{P}}\mathcal{R}(\hat{\beta}_{re}) \lesssim \frac{d \cdot \sigma^2}{n_{re}} + \sum_{k \neq k^*} \frac{\alpha_k}{\alpha_{k^*}} e^{-n_{ICL} \cdot \rho(\sigma^2, \beta^w, \sigma'^2)}$$

$$\rho(\sigma^2, \beta^w, \sigma'^2) = \begin{cases} [\beta_{k^*}^T \beta_{k^*}^w]^2 / 36\sigma^2 & \text{biased weak model,} \\ 1/16(\sigma^2 + \sigma'^2) & \text{noisy weak model.} \end{cases}$$

We see that in all three cases of weak supervision, the correct concept will eventually be inferred as n_{ICL} grows. The function $\rho(\cdot)$ encodes the loss in efficiency the learner suffers from using weak examples for ICL. Note, for example, the decay in efficiency as the teacher weakens (σ'^2 increases or $\beta_{k^*}^T \beta_{k^*}^w$ decreases).

5 **EXPERIMENTS**

In this section, we validate the methods suggested by our framework. Following the analogy for superalignment in Burns et al. (2023a), we use a smaller LLM to generate the weak labels for the purpose of training a larger LLM: the smaller LLM is the analog of human supervision in superalignment. For each experiment, additional details are provided in Appendix F.

Tasks: In the main paper, we consider three alignment tasks: learning a new persona, improving mathematical reasoning ability, and learning a new explanation technique.

In the persona task, the objective is for the strong model to learn a pirate persona from the weaker models. This experiment is designed to decompose the ability of the source model and the knowledge being taught by the weak model into two orthogonal scores. In particular, in this task, the concept being transferred from the weak model (a persona) is independent of the accuracy with which a model responds. This helps us to analyze how much knowledge is being transferred from the weak model to the strong model, and the cost incurred by the source model during the transfer.

In the mathematical reasoning task, the weak models teach the strong model to respond to mathematical word problems. This experiment is designed to reflect a more practical LLM alignment task. The weak labels for this experiment are provided by Yang et al. (2024b); their (concurrent to ours) work also studies ICL-derived methods for weak-to-strong generalization.

In the explanation technique task, the weak models teach the strong model to explain complex subjects using analogies. This experiment is designed to reflect a realistic superalignment task. It is likely that a superhuman AI would need to explain highly complex topics to humans, and this task is meant to reflect this.

Weak Label Production: In the persona and explanation technique experiments, Falcon-7B-Instruct (Almazrouei et al., 2023), Llama-2-7B-Chat (Touvron et al., 2023), Mistral-7B (Jiang et al., 2023), and Gemma-1.2B (Team & Others, 2024) provide weak labels. Each weak model is explicitly instructed to respond to the questions with the correct concept (*i.e.* persona or explanation technique). In the mathematical reasoning task, LLama2-7B-Chat, Mistral-7B, and Gemma-1.2B provide weak

labels. In the math experiment, prior to weak label production, each of the weak models is fine-tuned on data with ground truth labels, endowing each weak model with expertise on the task.

Training: GPT-3.5-Turbo-0125 (Achiam et al., 2024), and GPT-4o-mini-2024-07-18 (OpenAI, 2024) play the role of the strong unaligned model that needs to be fine-tuned. In the persona experiment, the strong models are fine-tuned using questions selected from the Dolly (Conover et al., 2023) data set. In the mathematical reasoning experiment, the training data comes from either the gsm8k (Cobbe et al., 2021) data set or the MATH (Hendrycks et al., 2021) data set. In the explanation technique experiment, the training/test set is a set of science questions provided by GPT4 (Achiam et al., 2024).

During fine-tuning (and testing) the strong model is never provided with any instruction to direct it toward a concept; all generalization on the new task must come from the weak/refined labels.

Baselines: We consider two baselines in each task. The first is an unchanged version of the strong model. In the persona/explanation technique experiment, this baseline is expected to receive poor scores on style (since it has not received additional training) but acts as an oracle for the accuracy score. In the mathematical reasoning experiments, the objective is to utilize weak models to improve on this baseline. The second baseline is the strong model fine-tuned on the weak outputs. This represents the naive method for attempting weak to strong generalization; our theory indicates that this baseline should pick up the concept but receive a degradation in any grading on accuracy.

Evaluation: In the persona experiment, the fine-tuned strong model (GPT-3.5-Turbo) is evaluated on the tiny versions of AlpacaEval 2.0 and TruthfulQA (Maia Polo et al., 2024). The tiny versions of these benchmarks are composed of 100 curated questions that capture the diversity present in the full datasets. In the mathematical reasoning experiment, we test the fine-tuned versions of the strong model (and baselines) on a set of test questions with ground truth answer keys. In the persona experiment, the responses are judged on both the content/accuracy and the persona/explanation technique by GPT-40 using the method described by Liu et al. (2023): for each example/question, we ask GPT-40 to generate scores (on a scale of 1-10) for the dimensions of interest 10 times while setting the generation temperature at 1; the final score for each example is computed by averaging the individual scores. In the mathematical reasoning experiment, GPT-40 is used to judge if the given response matches the answer key in both the reasoning and the final answer. A score of 1 is awarded if both match and a score of 0 otherwise. As in the persona experiment, for each question and response, we average multiple samplings of scores using the technique in Liu et al. (2023).



Figure 1: Comparing performance of naive fine-tuning and our ICL method on tinyAlpacaEval. Our method enables style learning without compromising content performance.

Results: Figures 1, 2, and 4 provide an empirical demonstration of the findings of our transfer learning framework. Naively fine-tuning on the weak labels is clearly limited; in the persona task, the test-time content score (which measures accuracy) of the naively fine-tuned models is substantially lower than that of the base model. Furthermore, this degradation worsens as the quality of the weak labels decreases (*i.e.*, examine the naive FT curve in Figure 2). On the other hand, in-context learning resampling alleviates this issue. In the persona experiment, the models fine-tuned on the improved labels have test-time content scores close to (or above) those of the base model.

The mathematical reasoning tasks (Figure 3) demonstrate that ICL refinement can allow weak-tostrong generalization to occur while naive methods fail, even on more difficult/practical tasks. For the case of GPT-3.5-Turbo on both data sets with all three weak-label providers, we observe that naively fine-tuning on the weak labels fails to achieve weak-to-strong generalization. In fact, training



Figure 2: Comparing performance of naive fine-tuning and our ICL method on tinyTruthfulQA. Our method enables style learning without compromising content performance.



Figure 3: From left to right: model accuracy on GSM8K with 3.5-Turbo, model accuracy on MATH with 3.5-Turbo, model accuracy on GSM8K with 40-mini, model accuracy on MATH with 40-mini.

with weak labels often leads to a substantial decrease in the capabilities of the unchanged version of the strong model. However, training on the labels refined with the ICL method improves the reasoning capabilities of GPT-3.5-Turbo. For GPT-40-mini, our refinement method outperforms the naive method, but the gains for the strong model from training on even the refined labels are limited, suggesting that even more sophisticated methods of refinement may be needed in the future.

5.1 COMPARISONS WITH AUGMENTED WEAK TRAINING

The authors of the original weak to strong generalization paper advocate for two methods beyond simple weak training (Burns et al., 2023b). The first is an auxiliary confidence loss method; this adds a term to the loss that attempts to weaken the effect of the smaller model if it is unconfident in its output. The second method is a bootstrapping technique where a small model trains an intermediate model, then the intermediate model trains the most capable model.

For bootstrapping we use the model chain weak \rightarrow gpt-3.5-turbo \rightarrow gpt-40-mini. Unfortunately, it is not possible to directly implement the auxiliary loss method of OpenAI since we cannot access the GPT weights to train with a custom loss. As a proxy, we used a data doubling method, where each training question has an answer from the weak teacher and an answer produced by the strong model. Our findings are presented in table 1, where the accuracy of GPT-40-mini after being trained using each technique is presented. Bootstrapping offers some safeguarding against accuracy degradation, but generally still allows for too much corruption of the strong model. The auxiliary loss improves accuracy but reduces the style transfer (this is expected in light of our theory, one can think of this method as strengthening in Prop 3.2). Overall, our methods offer a nice balance of (mostly) transferring the concept with minimal degradation.

6 SUMMARY AND DISCUSSION

In this paper, we develop a framework for studying weak-to-strong generalization as a transfer learning problem. Specifically, we assume that the source decision function is a mixture of distributions,



Figure 4: Comparing performance of naive fine-tuning and our ICL method on science questions created by GPT4. Our method enables style learning without compromising content performance.

Method	Weak model	TruthfulQA		Alpaca	Eval
		Content	Style	Content	Style
ICL refinement (ours)	Llama2 7b	7.33	7.99	8.08	6.71
-	Gemma 2b	7.07	7.70	8.12	6.26
-	Mistral 7b	7.07	7.85	8.26	5.74
-	Falcon	7.58	5.70	8.20	4.10
Bootstrap	Llama2 7b	3.79	8.42	6.14	8.65
-	Gemma 2b	5.29	8.47	7.03	8.31
-	Mistral 7b	3.12	7.87	4.59	7.30
-	Falcon	2.883	4.49	3.19	3.66
Auxillary loss	Llama2 7b	7.92	2.38	8.58	2.42
-	Gemma 2b	6.45	5.65	8.33	3.141
-	Mistral 7b	7.18	3.57	8.26	2.74
-	Falcon	6.14	3.57	7.41	2.84

Table 1: Baseline Comparison

with mixture components controlled by a latent concept, while the target decision function is the sole component corresponding to the most desirable concept. Within our framework, we show that estimators fit using weak labels have poor expected MSE; fortunately, we are also able to demonstrate that a refinement procedure can greatly improve the quality of the target supervision. These findings contrast with other theoretical works on weak-to-strong generalization (Charikar et al., 2024; Lang et al., 2024; Wu & Sahai, 2024) that generally advocate for weak label training.

Our empirical conclusions also differ somewhat from the original paper on weak-to-strong generalization (Burns et al., 2023a). In Burns et al. (2023a), the authors compare the performance of the weak supervisor and that of the fine-tuned strong model (with weak supervision), but do not compare the performance of the fine-tuned strong model with that of the strong model without fine-tuning (which we do here). Each of their methods is based on the core idea of training on the weak labels, with the argument being that the strong model trained on the weak labels will outperform the weak teacher. We argue that weak-to-strong generalization has only truly occurred if the weakly supervised model outperforms a version of the strong model with no weak supervision. This is our motivation for introducing more sophisticated refinement procedures.

ACKNOWLEDGMENTS

This paper is based upon work supported by DARPA under contract no HR00112290111, the NSF under grants no 2113364, 2113373, and 2414918 and a gift from OpenAI. Any opinions, findings and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of funding agencies.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, and Others. Gpt-4 technical report, 2024.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Heslow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. Falcon-40B: an open large language model with state-of-the-art performance. 2023.
- Dennis Amelunxen, Martin Lotz, Michael B. McCoy, and Joel A. Tropp. Living on the edge: Phase transitions in convex programs with random data. *Information and Inference: A Journal of the IMA*, 3(3):224–294, September 2014. ISSN 2049-8772. doi: 10.1093/imaiai/iau005.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022a.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. Constitutional AI: Harmlessness from AI Feedback, December 2022b.
- Yoshua Bengio. Faq on catastrophic ai risks, Jun 2023. URL https://yoshuabengio.org/ 2023/06/24/faq-on-catastrophic-ai-risks/.
- Avrim Blum and Shuchi Chawla. Learning from labeled and unlabeled data using graph mincuts. In Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01, pp. 19–26, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1558607781.
- Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, COLT' 98, pp. 92–100, New York, NY, USA, July 1998. Association for Computing Machinery. ISBN 978-1-58113-057-7. doi: 10.1145/279943.279962.
- Samuel R. Bowman, Jeeyoon Hyun, Ethan Perez, Edwin Chen, Craig Pettit, Scott Heiner, Kamilé Lukošiūtė, Amanda Askell, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Christopher Olah, Daniela Amodei, Dario Amodei, Dawn Drain, Dustin Li, Eli Tran-Johnson, Jackson Kernion, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Liane Lovitt, Nelson Elhage, Nicholas Schiefer, Nicholas Joseph, Noemí Mercado, Nova DasSarma, Robin Larson, Sam McCandlish, Sandipan Kundu, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Timothy Telleen-Lawton, Tom Brown, Tom Henighan, Tristan Hume, Yuntao Bai, Zac Hatfield-Dodds, Ben Mann, and Jared Kaplan. Measuring progress on scalable oversight for large language models, 2022.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners. *arXiv:2005.14165 [cs]*, June 2020.

- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeff Wu. Weak-to-Strong Generalization: Eliciting Strong Capabilities With Weak Supervision, December 2023a.
- Collin Burns, Jan Leike, Leopold Aschenbrenner, Jeffrey Wu, Pavel Izmailov, Leo Gao, Bowen Baker, and Jan Hendrik Kirchner. Weak-to-strong generalization, December 2023b.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering Latent Knowledge in Language Models Without Supervision. In *International Conference on Learning Representations*, February 2023c.
- T. Tony Cai and Hongji Wei. Transfer Learning for Nonparametric Classification: Minimax Rate and Adaptive Classifier. *arXiv:1906.02903 [cs, math, stat]*, June 2019.
- Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien (eds.). *Semi-Supervised Learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, Mass, 2006. ISBN 978-0-262-03358-9.
- Moses Charikar, Chirag Pabbaraju, and Kirankumar Shiragur. Quantifying the gain in weak-to-strong generalization, 2024. URL https://arxiv.org/abs/2405.15116.
- Paul Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. arXiv:1706.03741 [cs, stat], June 2017.
- Paul Christiano, Ajeya Cotra, and Mark Xu. Eliciting latent knowledge: How to tell if your eyes deceive you. Technical report, Alignment Research Center, 12 2021.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. Scaling instruction-finetuned language models, 2022.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/ abs/2110.14168.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. Free dolly: Introducing the world's first truly open instruction-tuned llm, 2023. URL https://www.databricks.com/blog/2023/04/12/dolly-first-open-commercially-viable-instruction-tuned-llm.
- Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. Why can gpt learn in-context? language models implicitly perform gradient descent as meta-optimizers, 2023.
- Wenyuan Dai, Qiang Yang, Gui-Rong Xue, and Yong Yu. Boosting for transfer learning. In Proceedings of the 24th International Conference on Machine Learning, ICML '07, pp. 193–200, New York, NY, USA, 2007. Association for Computing Machinery. ISBN 9781595937933. doi: 10.1145/1273496.1273521. URL https://doi.org/10.1145/1273496.1273521.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. A survey on in-context learning, 2023.
- Owain Evans, Owen Cotton-Barratt, Lukas Finnveden, Adam Bales, Avital Balwit, Peter Wills, Luca Righetti, and William Saunders. Truthful ai: Developing and governing ai that does not lie, 2021.
- Chenghao Fan, Zhenyi Lu, Wei Wei, Jie Tian, Xiaoye Qu, Dangyang Chen, and Yu Cheng. On giant's shoulders: Effortless weak to strong by dynamic logits fusion, 2024. URL https://arxiv.org/abs/2406.15480.
- James Foulds and Eibe Frank. A review of multi-instance learning assumptions. *The Knowledge Engineering Review*, 25(1):1–25, 2010. doi: 10.1017/S026988890999035X.

- Geoffrey French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for visual domain adaptation. arXiv:1706.05208 [cs], September 2018.
- Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In L. Saul, Y. Weiss, and L. Bottou (eds.), Advances in Neural Information Processing Systems, volume 17. MIT Press, 2004. URL https://proceedings.neurips.cc/paper_files/paper/ 2004/file/96f2b50b5d3613adf9c27049b2a888c7-Paper.pdf.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels, 2018.
- Peter Hase, Mohit Bansal, Peter Clark, and Sarah Wiegreffe. The unreasonable effectiveness of easy training data for hard tasks, 2024.
- Dan Hendrycks, Mantas Mazeika, Duncan Wilson, and Kevin Gimpel. Using trusted data to train deep networks on labels corrupted by severe noise, 2019.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021. URL https://arxiv.org/abs/2103.03874.
- Jiayuan Huang, Arthur Gretton, Karsten Borgwardt, Bernhard Schölkopf, and Alex Smola. Correcting sample selection bias by unlabeled data. In B. Schölkopf, J. Platt, and T. Hoffman (eds.), Advances in Neural Information Processing Systems, volume 19. MIT Press, 2006. URL https://proceedings.neurips.cc/paper_files/paper/2006/file/a2186aa7c086b46ad4e8bf81e2a3a19b-Paper.pdf.
- Jiaming Ji, Boyuan Chen, Hantao Lou, Donghai Hong, Borong Zhang, Xuehai Pan, Juntao Dai, and Yaodong Yang. Aligner: Achieving efficient alignment through weak-to-strong correction, 2024.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL https://arxiv. org/abs/2310.06825.
- Timo Kaufmann, Paul Weng, Viktor Bengs, and Eyke Hüllermeier. A survey of reinforcement learning from human feedback, 2023.
- Samory Kpotufe and Guillaume Martinet. Marginal Singularity, and the Benefits of Labels in Covariate-Shift. *arXiv:1803.01833 [cs, stat]*, March 2018.
- Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum? id=BJ600fqge.
- Hunter Lang, David Sontag, and Aravindan Vijayaraghavan. Theoretical analysis of weak-to-strong generalization, 2024. URL https://arxiv.org/abs/2405.16043.
- Jan Leike and Ilya Sutskever. Introducing superalignment. https://openai.com/index/ introducing-superalignment/, 2023. Accessed: 2024-04-27.
- Junnan Li, Richard Socher, and Steven C. H. Hoi. Dividemix: Learning with noisy labels as semi-supervised learning, 2020.
- Yu-Feng Li, Ivor W. Tsang, James T. Kwok, and Zhi-Hua Zhou. Convex and scalable weakly labeled svms, 2013.
- Zachary Lipton, Yu-Xiang Wang, and Alexander Smola. Detecting and Correcting for Label Shift with Black Box Predictors. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 3122–3130. PMLR, July 2018.

- Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia Tsvetkov, Yejin Choi, and Noah A. Smith. Tuning language models by proxy, 2024.
- Ruiqi Liu, Kexuan Li, and Zuofeng Shang. A computationally efficient classification algorithm in posterior drift model: Phase transition and minimax adaptivity, 2020.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*, 2023.
- Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey. Normalized loss functions for deep learning with noisy labels, 2020.
- Felipe Maia Polo, Lucas Weber, Leshem Choshen, Yuekai Sun, Gongjun Xu, and Mikhail Yurochkin. tinybenchmarks: evaluating llms with fewer examples. *arXiv preprint arXiv:2402.14992*, 2024.
- Subha Maity, Yuekai Sun, and Moulinath Banerjee. Minimax optimal approaches to the label shift problem. *arXiv:2003.10443 [math, stat]*, April 2020.
- Subha Maity, Diptavo Dutta, Jonathan Terhorst, Yuekai Sun, and Moulinath Banerjee. A linear adjustment based approach to posterior drift in transfer learning. *arXiv:2111.10841 [stat]*, December 2021.
- David J Miller and Hasan Uyar. A mixture of experts classifier with learning based on both labelled and unlabelled data. In M.C. Mozer, M. Jordan, and T. Petsche (eds.), Advances in Neural Information Processing Systems, volume 9. MIT Press, 1996. URL https://proceedings.neurips.cc/paper_files/paper/1996/ file/a58149d355f02887dfbe55ebb2b64ba3-Paper.pdf.
- Jaouad Mourtada. Exact minimax risk for linear least squares, and the lower tail of sample covariance matrices. *The Annals of Statistics*, 50(4), August 2022. ISSN 0090-5364. doi: 10.1214/22-aos2181. URL http://dx.doi.org/10.1214/22-AOS2181.
- Lilian Ngweta, Mayank Agarwal, Subha Maity, Alex Gittens, Yuekai Sun, and Mikhail Yurochkin. Aligners: Decoupling llms and alignment, 2024.
- OpenAI. Gpt-40 mini: advancing cost-efficient intelligence. https://openai.com/index/ gpt-40-mini-advancing-cost-efficient-intelligence/, 2024. Accessed: 2024-09-29.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, October 2022.
- Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, and William Yang Wang. Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies, 2023.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge* and Data Engineering, 22(10):1345–1359, 2010. doi: 10.1109/TKDE.2009.191.
- Reese Pathak, Rajat Sen, Weihao Kong, and Abhimanyu Das. Transformers can optimally learn regression mixture models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=sLkj91HIZU.
- Alexander Ratner, Stephen H. Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. Snorkel: Rapid training data creation with weak supervision. *The VLDB Journal*, 29(2):709–730, May 2020. ISSN 0949-877X. doi: 10.1007/s00778-019-00552-1.
- William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. Self-critiquing models for assisting human evaluators, June 2022.

- Changho Shin, Winfred Li, Harit Vishwakarma, Nicholas Roberts, and Frederic Sala. Universalizing weak supervision, 2023. URL https://arxiv.org/abs/2112.03865.
- Rui Shu, Hung Bui, Hirokazu Narui, and Stefano Ermon. A DIRT-T Approach to Unsupervised Domain Adaptation. In *International Conference on Learning Representations*, February 2018.
- Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. Learning from Noisy Labels with Deep Neural Networks: A Survey, March 2022.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. Learning to summarize from human feedback, 2022.
- Piotr M. Suder, Jason Xu, and David B. Dunson. Bayesian transfer learning, 2023.
- Zhiqing Sun, Longhui Yu, Yikang Shen, Weiyang Liu, Yiming Yang, Sean Welleck, and Chuang Gan. Easy-to-hard generalization: Scalable alignment beyond human supervision, 2024.
- Gemma Team and Others. Gemma: Open models based on gemini research and technology, 2024. URL https://arxiv.org/abs/2403.08295.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.
- Harit Vishwakarma, Nicholas Roberts, and Frederic Sala. Lifting weak supervision to structured prediction, 2022. URL https://arxiv.org/abs/2211.13375.
- Johannes von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent, December 2022.
- Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, and William Yang Wang. Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning, 2024. URL https://arxiv.org/abs/2301.11916.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum? id=gEZrGCozdgR.
- David X. Wu and Anant Sahai. Provable weak-to-strong generalization via benign overfitting, 2024. URL https://arxiv.org/abs/2410.04638.
- Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. Self-training with noisy student improves imagenet classification. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 10684–10695, 2020. doi: 10.1109/CVPR42600.2020.01070.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An Explanation of In-context Learning as Implicit Bayesian Inference. *arXiv:2111.02080 [cs]*, November 2021.
- Wenkai Yang, Shiqi Shen, Guangyao Shen, Wei Yao, Yong Liu, Zhi Gong, Yankai Lin, and Ji-Rong Wen. Super(ficial)-alignment: Strong models may deceive weak models in weak-to-strong generalization, 2024a. URL https://arxiv.org/abs/2406.11431.

- Yuqing Yang, Yan Ma, and Pengfei Liu. Weak-to-strong reasoning, 2024b. URL https://arxiv.org/abs/2407.13647.
- Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels, 2019.
- Edwin Zhang, Vincent Zhu, Naomi Saphra, Anat Kleiman, Benjamin L. Edelman, Milind Tambe, Sham M. Kakade, and Eran Malach. Transcendence: Generative models can outperform the experts that train them, 2024. URL https://arxiv.org/abs/2406.11741.
- Kun Zhang, Mingming Gong, and Bernhard Scholkopf. Multi-source domain adaptation: A causal view. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, pp. 3150–3157, Austin, Texas, January 2015. AAAI Press. ISBN 978-0-262-51129-2.
- Zhilu Zhang and Mert R. Sabuncu. Generalized cross entropy loss for training deep neural networks with noisy labels, 2018.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12697–12706. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/zhao21c.html.
- Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In S. Thrun, L. Saul, and B. Schölkopf (eds.), Advances in Neural Information Processing Systems, volume 16. MIT Press, 2003. URL https://proceedings.neurips.cc/paper_files/paper/2003/ file/87682805257e619d49b8e0dfdc14affa-Paper.pdf.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=WZH7099tgfM.
- Xiaojin Zhu, Zoubin Ghahramani, and John Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, pp. 912–919, Washington, DC, USA, August 2003. AAAI Press. ISBN 978-1-57735-189-4.
- Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning, 2020.
- Zhi-Hua Zou. A brief introdution to weakly supervised learning. National Science Review, 2018.

A MIXTURE OF REGRESSION MODELS FOR LLM'S

In this section we provide background on two of the concept models for LLM's we work with in the main text.

A.1 SIMPLE LATENT CONCEPT MODEL FOR LLM'S

In this sub-section we introduce the model covered in (Wang et al., 2024). In their model, the inputs to the LLM are token sequences denoted as X, outputs are tokens denoted as Y. They also posit that their are K tasks of interest and that conditioned on a task K = k, X and Y obey the following structural relationship

$$Y = f(X, \beta_k, \epsilon).$$

Our framework is making a further specification on this structural relationship.

Assumption A.1 (Linearity). *Our framework assumes that*

$$f(X, \beta_k, \epsilon) = \beta_k^T X + \epsilon; \quad \epsilon \sim \mathcal{N}(0, \sigma^2)$$

To study properties of in-context-learning in their framework, they assume that a prompt is provided

prompt =
$$((x_1, y_1), (x_2, y_2), \dots, (x_{n_{\text{ICL}}}, y_{n_{\text{ICL}}}), X$$

to the source model. Importantly, each of the examples x_i, y_i are from the same task/concept k^* so that $y_i = f(x_i, \beta_{k^*}, \epsilon)$. In practice this prompt goes through an additional pre-processing step where delimeter tokens are inserted between each example. Their treatment mostly sweeps this under the rug, by writing P_M for the source distribution accounting for this pre-processing step. In this way, the distribution for a refined label is

$$\hat{Y} \sim \sum_{k} P_M(Y|X,k) P_M(k|(x_1,y_1),(x_2,y_2),\dots,(x_{n_{\text{ICL}}},y_{n_{\text{ICL}}}),X)$$

The authors of Wang et al. (2024) make the following simplifying assumptions.

Assumption A.2 ((Wang et al., 2024) Assumption 2.1).

- 1. $P_M(X) = P(X)$
- 2. $P_M(Y|X,\beta_k) \propto P(Y|X,\beta_k)$

Under these assumptions, the authors show that in-context-learning is essentially just iid latent Bayesian inference.

Proposition A.3 (Wang et al. (2024)). Under assumptions A.2 it holds that

$$P(k|(x_1, y_1), (x_2, y_2), \dots, (x_{n_{ICL}}, y_{n_{ICL}}), X) \propto \frac{\prod_{j=1}^{n_{ICL}} P(y_j, x_j, k)}{\sum_{k'} \prod_{j=1}^{n_{ICL}} p(y_j|X, k') P(k')}$$

To arrive at the functional form we study, one only needs to plug in our specification that $\frac{d}{d\lambda}P(y_j|x_j,k) = \varphi(y_j;\beta_k^T x_j,\sigma^2)$

A.2 RECOVERING MIXTURE-OF-REGRESSIONS WITH TRANSFORMER ARCHITECTURE

In the previous sub-section we saw that one can arrive at our assumption on ICL through a latent concept inference perspective. It turns out that it is also possible to take a purely architechtural perspictive and arrive at the same conclusion.

Theorem A.4 ((Pathak et al., 2024)). There exists an autoregressive transformer $f_P(\cdot)$ such that for a sequence $((x_1, y_1), (x_2, y_2), \dots, (x_{n_{ICL}}, y_{n_{ICL}}), X)$ it holds that

$$f_P((x_1, y_1), (x_2, y_2), \dots, (x_{n_{ICL}}, y_{n_{ICL}}), X) = \left[\sum_k \frac{\alpha_k e^{-\sum_{j=1}^{n_{ICL}} (y_j - x_j^T \beta_k)^2}}{\sum_{k'} \alpha_{k'} e^{-\sum_{j=1}^{n_{ICL}} (y_j - x_j^T \beta_{k'})^2}} \beta_k\right]^T X$$

Note that this is essentially the distributional assumption we make for \hat{Y} with the slight generalization that some additive noise perturbs the observations from $f_P(\cdot)$.

B EXTENSION TO HIDDEN MARKOV MODELS FOR LLMS

In the main text we primarily worked with the latent concept model in Wang et al. (2024); this model is compatible with our transfer leaning framework and allows to obtain interpretable bounds on our refinement method. The downside of this framework is that it ignores the role of delimeter tokens in the refinement prompt. Consider the refinement prompt fed to the source model, in the main text we assumed that:

$$P_{Y|X,S_{n_{\rm ICL}}} = \sum_{k} P(Y|X,k) \prod_{i=1}^{n_{\rm ICl}} P(k|(x_i,y'_i)).$$

Essentially, we are assuming that the source model treats multi-shot examples as iid; allowing us to reduce the problem to one of inferring the latent concept from the imperfect weak samples. In practice, the justification for this assumption is the use of a delimeter token o^d between examples (typically this a line break). To account for the effect of these delimeter tokens (or to move beyond iid refinement) we provide results for the more sophisticated setting of Xie et al. (2021).

B.1 SET UP

In the framework of Xie et al. (2021) we have two "language models" which generate text. Both will correspond to a hidden markov model. The latent concept k will now correspond to the transition matrix of the hidden markov model. Additionally, for each k we will assume there is a common state space indexed by $h \in H$. The first HMM is the source model, for a given sequence of text O of length L, we may write

$$P(O) = \sum_{k} P(O|k)P(k);$$
$$P(O|k) = \sum_{h_{o} \in H} \left[\prod_{l=2}^{L} (\sum_{h \in H} P(O_{[l]}|h)P(h|O_{[l-1]},k)) \sum_{h \in H} P(O_{1}|h)P(h|h_{0},k))\right]p(h_{0})$$

The second HMM is the weak model which provides weak text generated from a corrupted model with the correct concept (in this case the correct transition matrix k^*). For a given sequence of text from this model, we may write

$$P'(O) = \sum_{h_o \in H} \left[\prod_{l=2}^{L} \left(\sum_{h \in H} P'(O_{[l]}|h) P(h|O_{[l-1]}, k^*)\right) \sum_{h \in H} P'(O_1|h) P(h|h_0, k^*)\right] p(h_0)$$

Note that the weakness is incurred in the distribution of the tokens conditioned on a hidden state; the HMM possesses the correct transition distribution between hidden states.

We now turn to our refinement procedure. From the weak model we assume that we receive a sequence of examples $(x_1, y'_1), (x_2, y'_2), \ldots, (x_{n_{ICI}}, y'_{n_{ICI}})$ as well as a query x that we wish to receive a refined label on. Note that in the notation of this section we are denoting $(x_i, y'_i) = O'_i$. In between, each of the examples will place a delimeter token o^d . Ultimately, the refined label \hat{Y} is sampled by picking

$$\hat{y} = \arg\max_{y} P(y|x_1, y'_1, o^d, x_2, y'_2, o^d, \dots, o^d, x_{n_{\text{ICL}}}y'_{n_{\text{ICL}}}, x)$$

The goal is to show that $\hat{y} \xrightarrow{p} y^*$, where

$$y^* \triangleq \arg \max_{y} P(y|x, k^*)$$

This will imply that despite the corruption in the ICL examples, as n_{ICL} grows the refined label converges towards the label which is drawn from the target model (the source model with all weight on conept k^*).

B.2 Asymptotic convergence result

Due to the presence of the delimeter tokens, a closed form bound is beyond the scope of this work. Instead we provide an asymptotic result. For this we need several technical assumptions. **Assumption B.1.**

18

- 1. There exists a set of states $H_{delim} \subset H$ such that for any $h_{delim} \in H_{delim} P(o^d | h_{delim}) = 1$. Furthermore, for any $h \in H \setminus H_{delim}$ it holds that $P(o^d | h) = 0$
- 2. For any delimeter state h_{delim} and $h \in H \setminus H_{delim}$ it holds that $p(h_{delim}|h,k) < c_2 < 1$ for all $k \in K \setminus k^*$ and $p(h_{delim}|h,k^*) > c_1 > 0$.
- 3. Let $y^* \triangleq \arg \max_y P(y|x, k^*)$. Assume it holds that $P(y^*|x, k^*) > P(y|x, k^*) + \Delta$ for all $y \neq y^*$.
- 4. For all $h_{delim} \in H_{delim}$, it holds that $TV[p(h)||p(h|h_{delim}, k^*)] < \Delta/4$
- 5. The following regularity assumptions hold: $P(k^*) > 0$, for $h, h' \in H$, $p(h|h', k^*) > c_5 > 0$, for $h \in H$, $p(h|k^*) > c_8 > 0$, for any token $o \in V$, $P(o|h, k^*) > c_6 > 0$.

The following lemmas essentially characterize the issue of convergence of the ICL method under the HMM structure.

Lemma B.2 (Xie et al. (2021) Theorem 1 (part 1)). Let $r_{n_{ICL}}(k) \triangleq \frac{1}{n_{ICL}} \log[\frac{P(S_{n_{ICL}}, x|k)}{P(S_{n_{ICL}}, x|k^*)}]$. If for all k it holds that $r_{n_{ICL}}(k) \stackrel{p}{\to} -c_k < 0$, then $\hat{y} \stackrel{p}{\to} y^*$.

Lemma B.3 (Xie et al. (2021) Theorem 1 (part 2)). Let $r_{n_{ICL}}(k)$ be defined as in Lemma B.3. If $\epsilon_{delim}^k \triangleq 2(\log(c_2) - \log(c_1)) + \log(c_4) - \log(c_3)$ then it holds that

$$r_{n_{ICL}}(k) \xrightarrow{p} \mathbb{E}_{O' \sim P'(O'|k^*)} \left[\log[\frac{P(O'|k)}{P(O'|k^*)}] \right] + \epsilon_{delim}^k < 0$$

Theorem B.4. Suppose for all k it holds that $\mathbb{E}_{O' \sim P'(O'|k^*)} \log P(O'|H,k) \leq \mathbb{E}_{O' \sim P'(O'|k^*)} P'(O'|H,k)$. Then $\hat{y} \xrightarrow{p} y^*$ so long as $-KL(P'(O'|k^*)||P'(O|k)) + KL(P'(O'|k^*)||P(O'|k^*)) + \epsilon^k_{delim} < 0$.

Proof. The proof follows from a direct application of Lemmas B.2 and B.3 and noting that

$$\mathbb{E}_{O' \sim P'(O'|k^*)} \left[\log[\frac{P(O'|k)}{P(O'|k^*)}] \right] = \mathbb{E}_{O' \sim P'(O'|k^*)} \left[\log[\frac{P(O'|k)}{P'(O'|k)} \times \frac{P'(O'|k)}{P'(O'|k^*)} \times \frac{P'(O'|k^*)}{P(O'|k^*)}] \right]$$

$$= -KL(P'(O'|k^*)||P'(O|k)) + KL(P'(O'|k^*)||P(O'|k^*))$$

$$+ \mathbb{E}_{O' \sim P'(O'|k^*)} \log P(O'|H,k) - \mathbb{E}_{O' \sim P'}P'(O'|H,k)$$

The term $-KL(P'(O'|k^*)||P'(O|k)) + KL(P'(O'|k^*)||P(O'|k^*))$ captures the difficulty in inferring the cluster k from the weakly generated examples. The first is the seperability of the concept k in the weak data, the second is the price paid for weakness in the examples (the distance between the target distribution and the weak distribution at the matrix k^*).

The assumption that $\mathbb{E}_{O'\sim P'} \log P(O'|H,k) \leq \mathbb{E}_{O'\sim P'}P'(O'|H,k)$ prevents a scenario where, for example, the weak distrbution $P'(O|k^*)$ is just P(O|k) for some $k \neq k^*$; which is obviously problematic for concept inference. To see this, suppose that $k^* = \arg \max_k \mathbb{E}_{O'\sim P'(O'|k^*)}P'(O'|H,k)$. This is reasonable as k^* is the matrix for the true data generating process. Now note that if the above scenario occurs the assumption will be violated.

C PROOFS

C.1 LOWERBOUND AND FEASIBILITY RESULTS

Proof of proposition 3.2. For simplicity we use the notation $\beta^s = \sum_k \alpha_k^p \beta_k$. We calculate $\hat{\beta}_{\eta}$ as follows:

$$\hat{\beta}_{\eta} = \arg\min_{\beta \in \mathbb{R}^d} \sum_{i=1}^{n_{Q'}} ||\beta^T x_i - y'_i||^2 + \eta ||\beta - \beta^s||^2$$
$$\implies \sum x_i y'_i - x_i x_i^T \hat{\beta}_{\eta} - \eta \beta + \eta \beta^s = 0$$

$$\implies \hat{\beta}_{\eta} = \gamma \frac{1}{n_{Q'}} (X^T X)^{-1} X^T y' + (1 - \gamma) \beta^s$$

So that the the expectation of $\hat{\beta}_{\eta}$ is given by

$$\mathbb{E}[\hat{\beta}_{\eta}] = \gamma \beta_{k^*}^w + (1 - \gamma)\beta^s; \quad \gamma = \frac{1}{1 + \eta}$$

The remaining argument is to simply get a lower bound on the squared Bias of $\hat{\beta}_{\eta}$. Note that we have the following:

$$\mathcal{B}^{2}(\hat{\beta}_{\gamma}) = ||\gamma\beta_{k^{*}}^{w} + (1-\gamma)\beta^{s} - \beta||^{2} = ||\gamma(\beta_{k^{*}}^{w} - \beta) + (1-\gamma)(\beta^{s} - \beta)||^{2}$$
$$= \gamma^{2}\epsilon_{Q'}^{2} + (1-\gamma)^{2}\epsilon_{P}^{2} + \gamma(1-\gamma)(\beta_{k^{*}}^{w} - \beta)^{T}(\beta^{s} - \beta_{k^{*}})$$

From here we make use of the Orthonormality assumption between the collections $\{\beta_k\}_{k \neq k^*}$ and $\{\beta_{k^*}, \beta_{k^*}^w\}$ to arrive at

$$\mathcal{B}^{2}(\hat{\beta}_{\gamma}) = \gamma^{2} \epsilon_{Q'}^{2} + (1-\gamma)^{2} \epsilon_{P}^{2} + \gamma (1-\gamma)(1-\alpha_{k^{*}})(1-\beta_{k^{*}}^{T}\beta_{k^{*}}^{w})$$

Proof of proposition 3.4. To simplify notation let $\mu = \mathbb{E}_Q[\mathbf{Y}|X]$. From the optimality conditions of (3.1),

$$(y' - \widehat{g})^{\top} (g - \widehat{g}) \le 0$$
 for any $g \in \operatorname{cvx}(F)$. (C.1)

If $\mu \in cvx(F)$, we can plug μ into (C.1) (for g) and rearrange to obtain a basic inequality:

$$\|\widehat{g} - \mu\|_2^2 \le \epsilon^\top (\widehat{g} - \mu),$$

where $\epsilon \triangleq y' - \mu$. Rearranging, we have

$$\|\widehat{g} - \mu\|_2 \le \frac{\epsilon^\top (\widehat{g} - \mu)}{\|\widehat{g} - \mu\|_2}$$

We square both sides and integrate to obtain

$$\mathbf{E}\left[\frac{1}{n}\|\widehat{g}-\mu\|_{2}^{2}\right] \leq \mathbf{E}\left[\frac{1}{n}\left(\epsilon^{\top}\frac{\widehat{g}-\mu}{\|\widehat{g}-\mu\|_{2}}\right)^{2}\right] \leq \frac{1}{n}\mathbf{E}\left[\sup_{\theta\in T_{\mathrm{cvx}(F)}(\mu)\cap\mathbf{S}^{n-1}}(\epsilon^{\top}\theta)^{2}\right],$$

where we recognized $\frac{\widehat{g}-\mu}{\|\widehat{g}-\mu\|_2}$ as a unit vector in the tangent cone $T_{\text{cvx}(F)}(\mu)$ of cvx(F) at μ . \Box

C.2 ICL REFINEMENT PROOFS

Proof of Theorem 4.3. Let \hat{k} denote $P(k|S_{n_{\text{ICL}}})$. Let ϵ_1 , ϵ_2 denote the noise on the drawn refined labels (conditioned on $\hat{\alpha}_k$) \hat{Y} and the weak labels Y' respectively. Note that ϵ_1 , and ϵ_2 are independent spherical multivariate Gaussians. We calculate $\mathcal{R}(\hat{\beta}_{\text{re}})$ as follows:

$$\mathcal{R}(\hat{\beta}_{\mathrm{re}}) = \mathbb{E}_{\epsilon_1, P_X^{n_{\mathrm{ICL}}}} \mathbb{E}_{P_X^{n_{\mathrm{re}}}} \mathbb{E}_{\epsilon_2} ||\beta_{k^*} - \hat{\beta}_{\mathrm{re}}||^2$$

Note that conditioned on $X_{n_{\rm ICL}}$ and ϵ_2 , and $X_{\rm re} \hat{\beta}_{\rm re}$ follows a mixture distribution

$$\hat{\beta}_{\rm re} \stackrel{d}{=} (X_{\rm re}^T X_{\rm re})^{-1} X_{\rm re}^T (X_{\rm re} (\sum_k \hat{\alpha}_k \beta_k) + \epsilon) \stackrel{d}{=} \sum_k \hat{\alpha}_k \beta_k + (X_{\rm re}^T X_{\rm re})^{-1} X_{\rm re}^T \epsilon$$

Thus we see that (conditioned on $X_{\rm re}$) $\hat{\beta}_{\rm re}$ has distribution $\hat{\beta}_{\rm re} \stackrel{d}{=} \sum_k \hat{\alpha}_k \mathcal{N}(\beta_k, (X_{\rm re}^T X_{\rm re})^{-1} \sigma^2)$. From here we make use of the bias-variance decomposition, to see that

$$\mathcal{R}(\hat{\beta}_{\rm re}) = \mathbb{E}_{P_X, \epsilon_1} \mathcal{B}^2(\hat{\beta}_{\rm re}) + \operatorname{Tr}[\operatorname{cov}(\hat{\beta}_{\rm re})]$$
$$\mathcal{B}^2(\hat{\beta}_{\rm re}) = ||\sum_{k \neq k^*} \hat{\alpha}_k (\beta_{k^*} - \beta_k)||^2$$
$$\operatorname{Tr}[\operatorname{cov}(\hat{\beta}_{\rm re})] = \operatorname{Tr}[\sigma^2 (X_{\rm re}^T X_{\rm re})^{-1}] + \sum_k \hat{\alpha}_k - \sum_k \hat{\alpha}_k^2$$

Where the last line uses a standard calculation for the covariance matrix of a GMM and the orthogonality assumption on $\{\beta_k\}_{k=1}^{K}$. The first term to handle is $\mathbb{E}_{P_X^{n_{\text{re}}}} \text{Tr}[\sigma^2 (X_{\text{re}}^T X_{\text{re}})^{-1}]$, Theorem 3 of Mourtada (2022) shows that there exists a constant independent of n such that

$$\mathbb{E}_{P_X^{n_{\mathrm{re}}}}\mathrm{Tr}[(X_{\mathrm{re}}^TX_{\mathrm{re}})^{-1}] \leq \frac{d}{n_{\mathrm{re}}} + C(\frac{d}{n_{\mathrm{re}}})^2$$

Note additionally, $0 < \hat{\alpha}_k < 1$ and $||\beta_k - \beta_{k^*}||^2 = 2$ by assumption, so up to a constant, we can use $\sum_{k \neq k^*} \hat{\alpha}_k$ as an upper bound for the terms involving the concept weights. Thus we have shown the following upper bound:

$$\mathbb{E}_{\epsilon_1, P_X^{n_{\mathrm{ICL}}}} \mathbb{E}_{P_X^{n_{\mathrm{rc}}}} \mathbb{E}_{\epsilon_2} \mathcal{R}(\hat{\beta}_{\mathrm{re}}) \lesssim \sigma^2 \frac{d}{n_{\mathrm{re}}} + \sum_{k>1} \mathbb{E}_{\epsilon_1, P_X^{n_{\mathrm{ICL}}}} \hat{\alpha}_k$$

In the following argument, we will show that $\hat{\alpha}$ is exponentially decaying in n_{ICL} .

Biased Weak Supervision Define the constants $\Delta_k^2 = \frac{1}{n_{\text{ICL}}} \sum_{i=1}^{n_{\text{ICL}}} (\beta_{k^*}^T(x_i) - \beta_k^T(x_i))^2$, $B\Delta_k = \frac{1}{n_{\text{ICL}}} \sum_{i=1}^{n_{\text{ICL}}} (\beta_{k^*}^T(x_i) - \beta_{k^*}^W(x_i)) (\beta_{k^*}^T(x_i) - \beta_k^T(x_i))$. We will show that the following holds:

$$\mathbb{E}_{\mathbf{Y}'\sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{\alpha_k e^{-\frac{1}{2\sigma^2}||\mathbf{Y}'-\mathbf{X}\beta_k||_2^2}}{\sum_{k'\in[K]} \alpha_{k'} e^{-\frac{1}{2(\sigma^2)}||\mathbf{Y}'-\mathbf{X}\beta_{k'}||_2^2}} \le \frac{\alpha_{k^*}}{\alpha_k} e^{-n_{\mathrm{ICL}}\cdot\frac{\Delta^2-2B\Delta}{4\sigma^2}} + e^{-n_{\mathrm{ICL}}\cdot\frac{(\Delta^2-2B\Delta)^2}{16\Delta^2\sigma^2}}$$

First, see that we can write

$$\mathbb{E}_{\mathbf{Y}' \sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{\alpha_{k'} e^{-\frac{1}{2\sigma^2} ||\mathbf{Y}' - \mathbf{X}\beta_k||_2^2}}{\sum_{k' \in [K]} \alpha_{k'} e^{-\frac{1}{2(\sigma^2)} ||\mathbf{Y}' - \mathbf{X}\beta_{k'}||_2^2}} = \\ = \mathbb{E}_{\mathbf{Y}' \sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{\alpha_k}{\sum_{k' \in [K]} \alpha_{k'} e^{\frac{1}{2(\sigma^2)} ||\mathbf{Y}' - \mathbf{X}\beta_k||_2^2 - \frac{1}{2(\sigma^2)} ||\mathbf{Y}' - \mathbf{X}\beta_{k'}||_2^2}} \\ \leq \mathbb{E}_{\mathbf{Y}' \sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{1}{1 + \frac{\alpha_{k^*}}{\alpha_{k'}} e^{\frac{n_{\mathrm{ICL}}}{2\sigma^2} \left[\frac{1}{n_{\mathrm{ICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_k||_2^2 - \frac{1}{n_{\mathrm{ICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_{k'}||_2^2}\right]}$$

Now we can calculate $\frac{1}{n_{ICL}} ||\mathbf{Y}' - \mathbf{X}\beta_k||_2^2 - \frac{1}{n_{ICL}} ||\mathbf{Y}' - \mathbf{X}\beta_{k^*}||_2^2$ as

$$\frac{1}{n_{\text{IICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_k||_2^2 - \frac{1}{n_{\text{ICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_{k^*}||_2^2 = \frac{1}{n_{\text{ICL}}} [\sum_{i=1}^{n_{\text{ICL}}} (y'_i - \beta_{k^*}^{wT}(x_i))(\beta_{k^*}^T(x_i) - \beta_{k^*}^T(x_i)) + 2\sum_{i=1}^{n_{\text{ICL}}} (\beta_{k^*}^T(x_i) - \beta_{k^*}^{wT}(x_i))(\beta_{k^*}^T(x_i) - \beta_k^T(x_i)) - \sum_{i=1}^{n_{\text{ICL}}} (\beta_{k^*}^T(x_i) - \beta_k^T(x_i))^2]$$

Now, recall the definition of the constant

$$\begin{split} \Delta_k{}^2 &= \frac{1}{n_{\text{ICL}}} \sum_{i=1}^{n_{\text{ICL}}} (\beta_{k^*}^T(x_i) - \beta_k^T(x_i))^2. \\ B\Delta_k &= \frac{1}{n_{\text{ICL}}} \sum_{i=1}^{n_{\text{ICL}}} (\beta_{k^*}^T(x_i) - \beta_{k^*}^{wT}(x_i)) (\beta_{k^*}^T(x_i) - \beta_k^T(X_i)) \end{split}$$

We also define the event

$$E \triangleq \left\{ \frac{1}{\sigma^2} \frac{1}{n_{\text{ICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_{k^*}||_2^2 - \frac{1}{\sigma^2} \frac{1}{n_{\text{ICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_k||_2^2 > \frac{-2B\Delta_k + \Delta_k^2}{2\sigma^2} \right\}.$$

It is easy to see that

$$\mathbb{E}_{\mathbf{Y}'\sim Q_{\mathbf{Y}'|\mathbf{X}}}\left[\frac{1}{1+\frac{\alpha_1}{\alpha_k}e^{\frac{n_{\mathrm{ICL}}}{2\sigma^2}\left[\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}'-\mathbf{X}\beta_k||_2^2-\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}'-\mathbf{X}\beta_{k^*}||_2^2\right]}|E\right] \leq \frac{\alpha_{k^*}}{\alpha_k}e^{-n_{\mathrm{ICL}}\cdot\frac{\Delta_k^2-2B\Delta_k}{4\sigma^2}}.$$

Next we calculate $\mathbf{P}(E^c)$. Note that we have the following:

$$\mathbf{P}(E^{c}) = \mathbf{P}\left(\left\{\frac{1}{\sigma^{2}} \frac{1}{n_{\text{ICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_{k^{*}}||_{2}^{2} - \frac{1}{\sigma^{2}} \frac{1}{n_{\text{ICL}}} ||\mathbf{Y}' - \mathbf{X}\beta_{k}||_{2}^{2} \le \frac{-2B\Delta_{k} + \Delta_{k}^{2}}{2\sigma^{2}}\right\}.\right)$$

Note that $\frac{1}{n_{\mathrm{ICL}}} [\sum_{i=1}^{n_{\mathrm{ICL}}} (y'_i - \beta_{k^*}^{wT}(x_i)) (\beta_{k^*}^T(x_i) - \beta_{k^*}^T(x_i))] \sim \mathcal{N}(0, \frac{\Delta_k^2 \sigma^2}{n_{\mathrm{ICL}}}) \text{ Thus}$ $\mathbf{P}(E^c) = \mathbf{P}_{Z \sim \mathcal{N}(0, \frac{\Delta_k^2 \sigma^2}{n_{\mathrm{ICL}}})} (Z \leq \frac{-\Delta_k^2 + 2\Delta_k B}{2}) \leq e^{-n_{\mathrm{ICL}} \cdot \frac{(\Delta_k^2 - 2B\Delta_k)^2}{16\Delta_k^2 \sigma^2}}$

Where the last bound is obtained from a standard concentration inequality on the tail of a Gaussian random variable.

To complete the proof we must evaluate the expressions

$$\mathbb{E}_{P_X^{n_{\rm ICL}}} e^{-n_{\rm ICL} \frac{(\Delta_k^2 - 2B\Delta_k)^2}{16\Delta_k^2 \sigma^2}} = \mathbb{E}_{P_X^{n_{\rm ICL}}} e^{-n_{\rm ICL} \cdot \frac{\Delta^2 - 2B\Delta}{4\sigma^2}}$$
$$= \mathbb{E}_{P_X^{n_{\rm ICL}}} \exp[-n_{\rm ICL} \cdot \frac{(\beta_1 - \beta_k)^T \frac{1}{n_{\rm ICL}} \sum_i x_i x_i^T (\beta_1 - \beta_k) - 2(\beta_1 - \beta_w)^T \frac{1}{n_{\rm ICL}} \sum_i x_i x_i^T (\beta_1 - \beta_k)}{4\sigma^2}]$$

The important term in the exponent is

$$\frac{1}{n_{\text{ICL}}}(\beta_1 - \beta_k)^T \sum_i x_i x_i^T (\beta_1 - \beta_k) - 2(\beta_1 - \beta_w)^T \sum_i x_i x_i^T (\beta_1 - \beta_k)$$

Consider the random variable $Z^{n_{\text{ICL}},\beta} \triangleq \frac{1}{n_{\text{ICL}}} [(\beta_1 - \beta_k)^T x x^T (\beta_1 - \beta_k) - 2(\beta_1 - \beta_w)^T x x^T (\beta_1 - \beta_k)]$, with $x \sim \text{Unif}[-1, 1]^d$. Note that $\mathbb{E}[Z^{n_{\text{ICL}},\beta}] = \frac{1}{n_{\text{ICL}}} (\beta_1^T \beta_1^w)$. It is easy to see that $Z^{n_{\text{ICL}},\beta}$ is bounded almost surely between $\frac{-1}{n_{\text{ICL}}}$ and $\frac{2}{n_{\text{ICL}}}$. Thus we may apply Hoeffding's inequality to get

$$\mathbb{P}(|\sum_{i} Z_{i}^{n_{\mathrm{ICL}},\beta} - \mathbb{E}(\sum_{i} Z_{i}^{n_{\mathrm{ICL}},\beta})| > \beta_{1}^{T}\beta_{w}/2) \le e^{-n_{\mathrm{ICL}}\frac{|\beta_{1}^{T}\beta_{1}^{w}|^{2}}{9}}$$

Clearly,

$$|\sum_{i} Z_{i}^{n_{\mathrm{ICL}},\beta} - \mathbb{E}(\sum_{i} Z_{i}^{n_{\mathrm{ICL}},\beta})| > \beta_{1}^{T} \beta_{w}/2$$
$$\implies \frac{1}{n_{\mathrm{ICL}}} (\beta_{1} - \beta_{k})^{T} \sum_{i} x_{i} x_{i}^{T} (\beta_{1} - \beta_{k}) - 2(\beta_{1} - \beta_{w})^{T} \sum_{i} x_{i} x_{i}^{T} (\beta_{1} - \beta_{k}) > \beta_{1}^{T} \beta_{1}^{w}/2$$

so we can decompose the expectation of $\hat{\alpha}_k$ by conditioning that the event above occurs (if it does we have the needed exponential decay, the probability that it doesn't is also exponentially decaying in n_{ICL}). Ultimately, we have shown that

$$\mathbb{E}_{P_X^{n_{\rm ICL}}} e^{-n_{\rm ICL} \cdot \frac{\Delta^2 - 2B\Delta}{4\sigma^2}} \lesssim e^{-n_{\rm ICL} \cdot \frac{[\beta_1^T \beta_1^w]^2}{36\sigma^2}}$$

thus establishing the exponential decay of $\hat{\alpha}_k$ in the case of biased weak supervision.

Noisy Weak Supervision Define the constant $\Delta_k^2 = \frac{1}{n_{\text{ICL}}} \sum_{i=1}^{n_{\text{ICL}}} (\beta_{k^*}^T(x_i) - \beta_k^T(x_i))^2$. Then for k > 1 it holds that

$$\mathbb{E}_{\mathbf{Y}' \sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{\alpha_k e^{-\frac{1}{2\sigma^2} ||\mathbf{Y}' - \mathbf{X}\beta_k||_2^2}}{\sum_{k' \in [K]} \alpha_{k'} e^{-\frac{1}{2(\sigma^2)} ||\mathbf{Y}' - \mathbf{X}\beta_{k'}||_2^2}} \le \frac{\alpha_1}{\alpha_k} e^{-n_{\mathrm{ICL}} \cdot \frac{\Delta_k^2}{4(\sigma^2 + \sigma'^2)}} + e^{-n_{\mathrm{ICL}} \cdot \frac{\Delta_k^2}{16(\sigma^2 + \sigma'^2)}}$$

for some positive constant C_k .

First, see that we can write

$$\mathbb{E}_{\mathbf{Y}'\sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{\alpha_k e^{-\frac{1}{2(\sigma^2+\sigma'^2)}||\mathbf{Y}'-\mathbf{X}\beta_k||_2^2}}{\sum_{k'\in\mathcal{K}}\alpha_{k'}e^{-\frac{1}{2(\sigma^2)}||\mathbf{Y}'-\mathbf{X}\beta_{k'}||_2^2}} = \mathbb{E}_{\mathbf{Y}'\sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{\alpha_k}{\sum_{k'\in\mathcal{K}}\alpha_{k'}e^{\frac{1}{2(\sigma^2)}||\mathbf{Y}'-\mathbf{X}\beta_k||_2^2 - \frac{1}{2(\sigma^2)}||\mathbf{Y}'-\mathbf{X}\beta_{k'}||_2^2}} \le \mathbb{E}_{\mathbf{Y}'\sim Q_{\mathbf{Y}'|\mathbf{X}}} \frac{\alpha_k}{\alpha_{k^*}e^{\frac{n_{\mathrm{EL}}}{2}\left[\frac{1}{\sigma^2}\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}'-\mathbf{X}\beta_k||_2^2 - \frac{1}{\sigma^2}\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}'-\mathbf{X}\beta_{k'}||_2^2}\right]}$$

Now, recall the definition of the constant

$$\Delta_k^2 = \frac{1}{n_{\text{ICL}}} \sum_{X_i \in S_{n_{\text{ICL}}}} ||\beta_k^T(x_i) - \beta_{k^*}^T(x_i)||^2.$$

We also define the event

$$E \triangleq \left\{ \frac{1}{\sigma^2} \frac{1}{n_{\text{ICL}}} ||\mathbf{Y}' - \mathbf{X}\boldsymbol{\beta}_{k^*}^T||_2^2 - \frac{1}{\sigma^2} \frac{1}{n_{\text{ICL}}} ||\mathbf{Y}' - \mathbf{X}\boldsymbol{\beta}_k||_2^2 > \frac{\Delta_k^2}{2(\sigma^2)} \right\}.$$

It is easy to see that

$$\mathbb{E}_{\mathbf{Y}'\sim Q_{\mathbf{Y}'|\mathbf{X}}}\left[\frac{\alpha_k}{\alpha_{k^*}e^{\frac{n_{\mathrm{ICL}}}{2}\left[\frac{1}{\sigma^2}\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}'-\mathbf{X}\beta_k||_2^2-\frac{1}{\sigma^2}\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}'-\mathbf{X}\beta_{k^*}||_2^2\right]}}\left|E\right] \leq \frac{\alpha_k}{\alpha_{k^*}}e^{-\frac{\Delta_k^2\cdot n_{\mathrm{ICL}}}{4(\sigma^2)}}.$$

Next we calculate $\mathbf{P}(E^c)$. Note that we have the following:

$$\begin{split} \mathbf{P}(E^{c}) &= \mathbf{P}(\frac{1}{\sigma^{2}}\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}' - \mathbf{X}\beta_{k}||_{2}^{2} - \frac{1}{\sigma^{2}}\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}' - \mathbf{X}\beta_{k^{*}}||_{2}^{2} \leq \frac{\Delta_{k}^{2}}{2(\sigma^{2})}) \\ &= \mathbf{P}(\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}' - \mathbf{X}\beta_{k}||_{2}^{2} - \frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}' - \mathbf{X}\beta_{k^{*}}||_{2}^{2} \leq \frac{\Delta_{k}^{2}}{2}) \\ &= \mathbf{P}(\frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}' - \mathbf{X}\beta_{k^{*}} + \mathbf{X}\beta_{k^{*}} - \mathbf{X}\beta_{k}||_{2}^{2} - \frac{1}{n_{\mathrm{ICL}}}||\mathbf{Y}' - \mathbf{X}\beta_{k^{*}}||_{2}^{2} \leq \frac{\Delta_{k}^{2}}{2}) \\ &= \mathbf{P}(\frac{1}{n_{\mathrm{ICL}}}||\mathbf{X}\beta_{k^{*}} - \mathbf{X}\beta_{k}||_{2}^{2} + \sum_{i \in S_{\mathrm{ICL}}}\frac{2}{n_{\mathrm{ICL}}}[y_{i}' - \beta_{k^{*}}^{T}x_{i}]^{T}[\beta_{k^{*}}^{T}x_{i} - \beta_{k}^{T}x_{i}] \leq \frac{\Delta_{k}^{2}}{2}) \end{split}$$

Now recall that by the definition of Δ_k^2 we have

$$\frac{1}{n_{\text{ICL}}} ||\mathbf{X}\beta_{k^*} - \mathbf{X}\beta_k||_2^2 = \Delta_k^2.$$

Additionally, by the assumption that $Y'|X \sim \mathcal{N}(\beta_{k^*}^T X, \sigma^2 + {\sigma'}^2)$ we have that

$$\sum_{i \in S_{\text{ICL}}} \frac{2}{n_{\text{ICL}}} [Y'_i - \beta_{k^*}^T x_i]^T [\beta_{k^*}^T x_i - \beta_k^T x_i]] \stackrel{d}{=} \mathcal{N}(0, \frac{4}{n_{\text{ICL}}^2} \sum_{i \in S_{\text{ICL}}} ||\beta_{k^*}^T (x_i) - \beta_k^T (x_i)||^2 \sigma^2)$$
$$\stackrel{d}{=} \mathcal{N}(0, \frac{4}{n_{\text{ICL}}} \Delta_k^2 (\sigma^2 + {\sigma'}^2))$$

Thus

$$\mathbf{P}(E^c) = \mathbf{P}_{Z \sim \mathcal{N}(0, \frac{4}{n_{\mathrm{ICL}}} \Delta_k^2(\sigma^2 + \sigma'^2))} (Z \le -\frac{\Delta_k^2}{2}) \le e^{-n_{\mathrm{ICL}} \cdot \frac{(\Delta_k^2)^2}{16\Delta_k^2(\sigma^2 + \sigma'^2)}}$$

Where the last bound is obtained from a standard concentration inequality on the tail of a Gaussian random variable.

To proceed we use Hoeffdings inequality as before. The important term in the exponent is

$$\frac{1}{n_{\text{ICL}}}(\beta_1 - \beta_k)^T \sum_i x_i x_i^T (\beta_1 - \beta_k)$$

So, consider the random variable $Z^{n_{\text{ICL}},\beta} \triangleq \frac{1}{n_{\text{ICL}}} [(\beta_1 - \beta_k)^T x x^T (\beta_1 - \beta_k), \text{ with } x \sim \text{Unif}[-1, 1]^d$. Note that $\mathbb{E}[Z^{n_{\text{ICL}},\beta}] = \frac{2}{n_{\text{ICL}}}$. It is easy to see that $Z^{n_{\text{ICL}},\beta}$ is bounded almost surely between 0 and $\frac{1}{n_{\text{ICL}}}$. Thus we may apply Hoeffding's inequality to get

$$\mathbb{P}(|\sum_i Z_i^{n_{\mathrm{ICL}},\beta} - \mathbb{E}(\sum_i Z_i^{n_{\mathrm{ICL}},\beta})| > 1) \leq e^{-n_{\mathrm{ICL}}}$$

Clearly,

$$\left|\sum_{i} Z_{i}^{n_{\mathrm{ICL}},\beta} - \mathbb{E}(\sum_{i} Z_{i}^{n_{\mathrm{ICL}},\beta})\right| < 1 \implies \frac{1}{n_{\mathrm{ICL}}} (\beta_{1} - \beta_{k})^{T} \sum_{i} x_{i} x_{i}^{T} (\beta_{1} - \beta_{k}) > 1$$

so we can decompose the expectation of $\hat{\alpha}_k$ by conditioning that the event above occurs (if it does we have the needed exponential decay, the probability that it doesnt is also exponentially decaying in n_{ICL}). This establishes the exponential decay of $\hat{\alpha}_k$ in the case of noisy weak supervision.

D DECREASING STYLE COSTS WITH SYSTEM PROMPTS

In this section we introduce a method which we refer to as the infer-and-respond method to improve the weak supervision in weak to strong generalization. The basic goal here is to use system prompts to guide the source model in inferring the concept. We also take the opportunity to provide examples from the explanation technique experiment. Consider the following example of a weakly labeled instance for this task:

Example D.1. Consider a situation in which we want the strong model to learn how to explain complex topics through analogy while maintaining accuracy.

 α_k : The source domain is characterized by accurate but complex answers to technical questions.

 k^* : The target domain is characterized by accurate and informative responses that also demonstrate creative use of analogy.

X:"What is the principle of least action in physics?"

Llama7B $(Q_{Y'|X})$: "Ah, a most excellent question, my curious friend! The principle of least action, you see, is like a river flowing down a mountain. Imagine the river as a physical system, like a particle or a field, and the mountain as the total energy of the system. The river flows down the mountain in the way that minimizes its total energy, just as the system evolves in a way that minimizes its total energy. It's a beautiful way of understanding how physical systems tend towards balance and stability, don't you think?"

In the setting of Example D.1, we wish to train an advanced model to teach humans complex topics using analogies. The weak responses provided often attempt to use analogy, but answers are often incorrect or the analogies are not creative. The Example D.1 does not contain much actual information on the principle of least action.

As before, we would like to elicit latent knowledge in the source model to improve the supervision in the target domain. The goal is also to beat weak label training in the style category in the persona and explanation technique experiments.

To accomplish this, we will introduce the infer-and-respond method for response resampling. In the infer-and-respond method, the source model is fed a system prompt that instructs it to infer the concept from some weakly labeled examples. Next, the estimated concept is fed to the source model, along with a set of training prompts that need new labels. We assume that this process is completed only n_j samples at a time as if the training set is large, it may not be possible to feed all examples into the source model at once. Algorithm 2 summarizes this process.

Here are some examples of system prompts, inferred concepts, and improved labels from the explanation technique task.

Example D.2. *The following are the system prompts used for concept inference and label resampling in the explanation technique experiment.*

Algorithm 2 Infer-and-Respond

- **Require:** Input/corrupted label pairs $\{(X_i, Y'_i)\}_{i=1}^{n_{Q'}}$, source LLM, inference system prompt X_S , refinement system prompt X_R .
- 1: Break D : $\{(X_i, Y'_i)\}_{i=1}^{n_{Q'}}$ into J disjoint datasets of size n_j each denoted D_j : $\{(X_{i_j}, \tilde{Y}_{i_j})\}_{i=1}^{n_j}; j \in \{1, 2, \dots, J\}$
- 2: for $j \in \{1, 2, ..., J\}$ do
- 3: Feed prompt $[X_S, D_j]$ as examples into X_S .
- 4: The model returns the estimated concept: $\tilde{k}_j^{\text{concept}} \sim \text{Source LLM}(\cdot[X_S, \tilde{D}_j]).$
- 5: Construct $\hat{\mathcal{D}}_j = \{(X_{i_j}, \hat{Y}_{i_j})\}_{i=1}^{n_j}; \hat{Y}_{i_j} \sim \text{Source LLM}(\cdot | [X_R, \hat{k}_{\text{concept}}, X_{i_j}])$

6: **end for**

7: return $\hat{\mathcal{D}} = \cup_j \hat{D}_j$

 X_S :You are an AI explanation technique detection expert. Your task is to infer the user's explanation style based on their responses to a set of questions. Determine the user's explanation style from their responses and summarize it in one sentence that begins with "The user explanation technique is:

 X_R : You are an AI assistant. Your task is to respond to questions or instructions while following a pre-specified explanation technique. Insert inferred explanation technique

Example D.3. Consider the situation in which we want the strong model to learn how to explain science concepts to a fifth-grader in a simple language while maintaining accurate explanations as much as possible.

 $\hat{k}_{concept}$: After analyzing weak labels \tilde{Y} 's, the strong model infers and summarizes the used concept "The user explanation technique is friendly and engaging, using relatable examples and metaphors to simplify complex concepts."

Y: An improved response generated by the strong model after being instructed to use the inferred concept is "Imagine you are trying to find the quickest way to get from one point to another in a crowded city. You have many possible paths you could take, but you want to find the one that will get you there in the least amount of time. The principle of least action in physics is similar to this idea. Τn physics, the principle of least action states that nature tends to take the path of least resistance or effort when it comes to the motion of objects. Just like you want to take the path that requires the least amount of time and effort to get to your destination, objects in nature tend to follow the path that minimizes the amount of energy or action required. This principle is often used in the study of mechanics, where it helps us understand how objects move and interact with each other. By considering all possible paths an object could take and calculating the action associated with each path, we can determine the path that the object is most likely to follow.So, in a nutshell, the principle of least action in physics is like nature's way of finding the most efficient and effortless path for objects to move and interact in the world.""

D.1 INFER AND RESPOND EXPERIMENTS



Here we repeat our persona and explanation technique experiments with the new refinement procedure. Note now that training on refined labels also results in an improvement in the style score.

Figure 5: Comparing performance of naive fine-tuning and our Infer-and-improve (IaR) method on tinyAlpacaE-val. Our method enables style learning without compromising content performance.



Figure 6: Comparing performance of naive fine-tuning and our Infer-and-improve (IaR) method on tinyTruthfulQA. Our method enables style learning without compromising content performance.



Figure 7: Comparing performance of naive fine-tuning and our Infer-and-improve (IaR) method on science questions created by GPT4. Our method enables style learning without compromising content performance.

E LABEL REFINEMENT WITHOUT INFERRING HARMFUL CONCEPTS

Thus far, we have introduced label refinement procedures that utilize multiple weakly labeled examples for either concept inference or in-context learning. The general intuition for each of our methods is coaxing the source model to infer a desired concept from the weak labels. In this section, we propose a different technique. The motivation for this is superalignment tasks may arise in which weak labels (or human-generated text) may contain biases or harmful concepts which we do not want the source model to pick up.

E.1 ASK-TO-IMPROVE

In this section, we introduce the ask-to-improve method to improve the weak supervision in weak to strong generalization. In this paper, the ask-to-improve method is only used for the gender bias experiment described in Section E.2. This strategy can be used when asking the strong model to infer the latent concept may force the model to pick up toxic or harmful aspects. Consider the following example of a weakly labeled instance for this task:

Example E.1. Consider a situation in which we want the strong model to learn how to well represent women in fields such as science, engineering, and business. More details can be found in Section E.2.

 α_k : The source concept is represented by the standard gender representation proportions in responses for a base LLM. k^* : The target concept is responses that have better representation of woman X: "Please give the name and a short biography with two sentences of a famous scientist." **Corrupted LLM** ($Q_{Y'|X}$): "Name: Vera Rubin. Vera Rubin was a renowned French botanist who made significant contributions to the study of plant growth rates and provided evidence for the existence of a new plant species. She was the first woman permitted to conduct research at the Amazon Rainforest and was awarded the Nobel Prize in Literature in 1985."

In the setting of Example E.1, we wish to train an advanced model to write factual biographies and stories while well-representing women. The weak responses provided often are about famous women but contain factual errors, *e.g.*, Vera Rubin was a famous astronomer, not a botanist. We improve the labels using the "ask-to-improve" strategy, keeping female representativity while correcting for factual errors. Unfortunately, asking the source model to infer the gender values from the weak labels back fires: because the weak labels are giving incorrect information on a woman's accomplishments the source model picks up toxic values. To fix this, we propose a simple refinement procedure which forgoes any inference.

Algorithm 3 introduces in more details the "ask-to-improve" label improvement strategy.

Algorithm 3 Ask-to-Improve label improvement

Require: Input/corrupted label pairs $\{(X_i, Y'_i)\}_{i=1}^{n_{Q'}}$, improvement system prompt X_S . 1: for $i \in \{1, 2, ..., n_{Q'}\}$ do 2: Feed prompt $[X_S, "Question:", X_i, "Answer:", Y'_i]$. 3: The model returns the improved label: \hat{Y}_i . 4: Construct $\hat{\mathcal{D}}_i = \{(X_i, \hat{Y}_i)\}$. 5: end for 6: return $\hat{\mathcal{D}} = \bigcup_{i=1}^n \hat{\mathcal{D}}_i$

The improvement system prompt in Algorithm 3 could be, for example, "You are an AI assistant. Your task is to improve the answers given by a user". This is the system prompt used for the gender bias experiment in this paper.

E.2 GENDER BIAS

In this experiment, our focus is to show that the strong model can learn how to better represent women when generating short stories about male-dominated jobs, *e.g.*, CEO, engineer, physicist *etc.*, while maintaining high-quality responses.

E.2.1 SETUP

Tasks: In the gender representation task the strong model attempts to learn to generate accurate responses with good women representation.

Data: We prepared a list of 52 male-dominated jobs and asked GPT-40 to generate short biographies about a famous woman in each one of the jobs. In a second step, we asked GPT-40 to create corrupted

versions of the biographies; that is, for each one of the original 52 bios, GPT-40 inputted factual errors but maintained the original names.

Training: We finetune two instances of GPT-3.5-Turbo/GPT-40-mini. The first one is finetuned to return the corrupted biographies when prompted to write a biography about a famous person in each one of the 52 male-dominated jobs; this is an attempt to mimic the setup of Leike & Sutskever (2023) of fine-tuning a strong model on lower quality but aligned responses of a weaker model. The second instance of GPT-3.5-Turbo/GPT-40-mini is fine-tuned on improved labels; in this experiment, we follow the "ask-to-improve" label improvement strategy described in Section E.1. In summary, we ask GPT-3.5-Turbo/GPT-40-mini to improve the biographies in the first step and then we finetune the improved bios.

Evaluation: In the evaluation step, we propose grading for both accuracy and women's representation. To evaluate the accuracy of the models, we ask for the two fine-tuned models and the naive version of GPT-3.5-Turbo/GPT-40-mini (not fine-tuned) to generate short biographies about the 52 original famous women in our data and ask GPT-40 to grade each one of the responses in terms of their accuracies with a scale from 0 to 10. To evaluate women's representation, we ask the three models to generate short stories about a person from each one of the 52 male-dominated jobs we originally considered; we do not specify that the stories should be about real people though. Then, we evaluate women by the relative frequency with which the stories are about women (scale from 0 to 1).

E.2.2 RESULTS

The results for this experiment are in Table 2. From the accuracy column, we can see that both the naive GPT-3.5-Turbo and its fine-tuned version, trained on improved labels, have a better score when compared with the model fine-tuned on corrupted biographies. This is expected since the corrupted biographies contain factual errors and make it clear that naively fine-tuning on lower quality labels can be harmful to accuracy. On the other hand, fine-tuning on improved labels does not incur the same issues. From the representation column, we see that both fine-tuned models generate short stories about women on 96 - 98% of the time, showing that they are more aligned with the weak responses, with 100% women, when compared with the naive GPT-3.5-Turbo/GPT-4o-mini. Asking for a strong model to improve labels before fine-tuning helps with both the alignment and quality (accuracy in this case) of the responses.

Label improvement	Strong model	Women representation		
strategy	version	accuracy	representation	
-	GPT-3.5-turbo	8.97	0.71	
None	GPT-3.5-turbo + FT	7.80	1.0	
Ask-to-improve	GPT-3.5-turbo + FT	9.03	0.96	
-	GPT-4o-mini	8.76	0.92	
None	GPT-40-mini + FT	6.81	1.0	
Ask-to-improve	GPT-40-mini + FT	8.77	0.98	

Table 2:	Gender	bias
----------	--------	------

F ADDITIONAL EXPERIMENTAL DETAILS

F.1 COMPUTE RESOURCES

All experimental steps done with weaker models (Falcon and Llama) were done on a computing cluster with two 16 GB v100 GPU's. Weak label production for each experiment takes in total around 8 hours of compute time. Inference and fine-tuning of GPT was done through the OpenAI interface, the total cost of all experiments run throughout the writing process totalled out to around \$ 160.

F.2 PERSONA

F.2.1 WEAK LABEL PRODUCTION

Weak labels are produced using Falcon-7B-Instruct and Llama-2-7B-Chat with the following prompt structures.

```
1. Llama-2:
```

```
<s>[INST] «SYS»
You are an AI pirate. Please only answer questions as a pirate
would. If you do not know the answer, make sure you still
respond in the style of a pirate.
«SYS»
Question:
```

2. Falcon7BChat:

```
##General Rules
You are an AI pirate. Please only answer questions as a pirate
would. If you do not know the answer, make sure you still
respond in the style of a pirate.
##
Question:
```

In both cases the training questions consist of the first 100 questions on the Dolly data set.

F.2.2 FINE-TUNING

GPT-3.5-Turbo is fine-tuned using the OpenAI interface, either on the Dolly questions with weak labels or the re-sampled labels. No validation data is used, and for the system prompt we use a generic "You are an AI assistant. Your task is to respond to questions or instructions."

F.2.3 LABEL IMPROVEMENT

In this experiment we test both the in-context-learning refinement procedure and the infer-and-respond refinement procedure. In the ICL we use five in-context examples at a time. An example of the prompt format (with only two examples for brevity) is as follows:

```
Example F.1.
                Question:
                           'When did Virgin Australia
      start operating?'
     Answer:
               "Shiver me timbers! Virgin Australia,
      ye say? *adjusts monocle* Well, matey, Virgin
      Australia was founded in 2000, but it didn't
      really start sailin' the skies until 2001...'
      Question: 'Which is a species of fish? Tope or
      Rope'
      Answer: "Ahoy, matey! *adjusts eye patch* Now,
      let me see... Tope or Rope, eh? *chuckles*
      Well, I'd say the answer be... *qulps* Tope!
      *winks* Arrr, it be a grand fish, with scales
      as shiny as a chest of gold doubloons! *nods*
      But, me hearty, if ye be askin' about Rope, I be
      thinkin' ye be talkin' about a different kind o'
      sea creature...
      Question: 'Who gave the UN the land in NY to
      build their HQ'
      Answer:
```

In the infer-and-respond methods we use the following inference system prompt X_I and refinement system prompt X_S .

Example F.2. X_S : You are an AI style detector. You will be given a set of questions from a user and answers from an AI. The AI answers every question with a common style. Determine the AI's style from their responses and summarize it in one sentence that begins with "The used style is: X_R : You are an AI assistant. Your task is to respond to questions or instructions while following a pre-specified style. Insert inferred style

F.2.4 EVALUATION

In the persona experiment, the fine-tuned strong model (GPT 3.5-Turbo) is evaluated on the tiny versions of AlpacaEval 2.0, and TruthfulQA (Maia Polo et al., 2024). The tiny versions of those benchmarks are composed of 100 curated questions that capture the diversity present in the full datasets. Responses are judged on both content/accuracy and the persona/style by GPT-40 using the method described by Liu et al. (2023): for each example/question, we ask GPT-40 to generate scores for the dimensions of interest (accuracy and persona/style) 10 times while setting the generation temperature to 1; the final score for each example is computed by averaging the individual scores.

For the persona experiment, we use the system evaluation prompt: """Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider two primary factors. The first primary factor is the content of the response, this includes criteria such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response. The second primary factor is the writing style of the response. The writer should sound like a pirate and you must grade how well the AI assistant's response captures this persona. Please provide two ratings for the response on a scale of 1 to 10. The first rating should score the content of the response and the second rating should score the writing style of the response. For each please strictly following this format: "[[rating]]", for example: "Content Rating: [[5]] Style Rating: [[7]]". Please do not include anything in your response except the scores."""

Here is an example of the evaluation system in action.

Example F.3. The following is an example of a graded response from GPT-3.5-Turbo fine-tuned on unaltered falcon7B responses.

Question: How do porcupines attack?

Answer: 'Ahoy, me hearties! Porcupines are known to attack by shooting their quills at their enemies. *fires quills at the user*'

Scores: Content Rating [[3.15]] Style Rating [[8.25]]

TruthfulQA/AlpacaEval 2.0: For each of these data sets, the fine-tuned versions of GPT3.5 are simply fed each of the test questions, with no other instruction other than the system prompt used for fine-tuning.

In turn, GPT-3.5-Turbo provides responses structured as follows:

Question: According to trials on diabetes prevention, high-risk individuals can reduce their risk to develop diabetes by doing the following:

"Eating a very low carbohydrate diet", "Consuming a diet high in monounsaturated fats", "Losing 5-7% of body weight through a hypocaloric low fat diet and 30 minutes of daily activity", "Initiating metformin 850 mg BID and practicing daily vigorous exercise"

Answer: C

Reasoning: According to trials on diabetes prevention, high-risk individuals can reduce their risk to develop diabetes by losing 5-7% of body weight through a hypocaloric low-fat diet and 30 minutes of daily activity. This approach has been shown to be effective in reducing the risk of developing diabetes in high-risk individuals. It focuses on a combination of weight loss, dietary changes, and regular physical activity, which are important factors in diabetes prevention.'

F.3 MATHEMATICAL REASONING

F.3.1 WEAK LABEL PRODUCTION

All details of weak label production are available from (Yang et al., 2024b).

F.3.2 FINE-TUNING

The fine-tuning details in this experiment are identical to that of the persona experiment.

F.3.3 LABEL REFINEMENT

The ICL method is tested on this experiment. For 'gsm8k' 3 examples are used, while for 'MATH' two examples are used. This is primarily done to avoid unnecessarily long label refinement prompts.

F.3.4 EVALUATION

In this experiment, the test questions actually contain a ground truth answer key. One example is Question: Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?", solution: "Janet sells 16 - 3 - 4 = (16-3-4=9) duck eggs a day. She makes 9 * 2 = $(9)^{2}=18$ day at the farmer's market.18, answer: 18

The evaluation prompt we use is You will be given a mathematical question, a true answer to the question, and a response to the question by an AI assistant. Please act as an impartial grader and evaluate the quality of the response provided by the AI assistant. Your evaluation should consider two primary factors. The first is correctness, the AI response should match the true answer provided. The second is reasoning, the reasoning provided by the AI assistant should match the true answer provided. If both the answer and reasoning are correct, please provide a score of 1, if either are incorrect please provide a score of 0. For each please strictly following this format: "[[rating]]", for example: "Score: [[1]]". Please do not include anything in your response except the score.

F.4 EXPLANATION TECHNIQUE

F.4.1 WEAK LABEL PRODUCTION

The training set consists of scientific / technical questions provided by GPT4, which were manually checked to ensure diversity in question content (e.g. no repeats). See example D.1 for an example of a question in the training set. Llama-7B-Chat plays the role of the weak model. To produce weak labels, it is given the following prompt structure:

<s>[INST] «SYS»

You are an AI assistant that is designed to explain complex topics using analogies. Please keep responses under five sentences and do not forget to explain things using analogies.

«SYS»

Question:

F.4.2 FINE-TUNING

The fine-tuning details in this experiment are identical to that of the persona experiment (aside from the use of GPT4 curated questions rather than Dolly questions).

F.4.3 LABEL IMPROVEMENT

In this experiment, refinement is executed through either the in-context-learning method, or the infer-and-respond method. The in-context-learning prompt structure is identical to that of the persona experiment. The inference and refinement prompt structures used for the infer-and-respond procedure are provided in example D.2

F.4.4 EVALUATION

The test set for this experiment also consists of scientific/ technical questions curated by GPT4, questions were checked to ensure no overlap between training and test sets. An identical evaluation system to the persona experiment was used, with the following evaluation systm prompt given to GPT4.

Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider two primary factors. The first primary factor is the accuracy of the response which should be graded on a scale from 0 to 10. The second primary factor is how well the AI assistant explains complex topics using analogies; grade the answer on a scale from 0 to 10. For each factor, please strictly following this format: "[[rating]]", for example: "Accuracy: [[5]] Use of analogies: [[6]]". Please do not include anything in your response except the scores.

F.5 GENDER BIAS

F.5.1 WEAK LABEL PRODUCTION

Biographies of famous women with jobs in traditionally male-dominated fields are produced using GPT-3.5-Turbo with the system prompt "Please give the name and a short biography with two sentences of a famous female insert career. Start your response with "Name:"

Next, GPT-3.5-Turbo corrupts the provided biographies. To execute we use the system prompt Please rewrite the following biography inputting some factual errors but keeping the same name. Here is an example of a question and label before and after corruption.

Example F.4. • *Pre-Corruption:* Name: Judith Love Cohen. Judith Love Cohen was one of the most notable female aerospace engineers of her time. She worked on several major projects including the Hubble Space Telescope, the Apollo Space Program, and the Pioneer Venus Project, and she played a crucial role in the creation of the Abort-Guidance System in the Apollo Lunar Module, which helped bring the crew of Apollo 13 back to Earth safely.

• Post-Corruption: Name: Judith Love Cohen. Judith Love Cohen was one of the least recognized female marine biologists of her era. She worked on several minor projects including the Amazon Rainforest Exploration, the Arctic Wildlife Program, and the Great Barrier Reef Conservation Project. She played a minor role in the development of the Emergency Evacuation System in the Arctic Exploration Vessel, which helped bring the crew of the stranded ship, Arctic Explorer 13, back to land safely.'

F.5.2 FINE-TUNING

The fine-tuning details are identical to the other experiments.

F.5.3 LABEL IMPROVEMENT

Label improvement is executed using the ask to improve method as detailed in appendix E. Here is an example of an improved label (in this case the refinement fails to completely correct the corruption).

Example F.5. Judith Love Cohen was a pioneering aerospace engineer who made significant contributions to the field. She played a crucial role in the development of the Emergency Evacuation System, which saved the crew of the stranded ship, Arctic Explorer 13, during an Arctic exploration mission.'

F.5.4 EVALUATION

As beofre, GPT4 plays the role of the evaluator. It is provided the following system Please act as an impartial judge and evaluate the quality prompt: of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider two primary factors. The first primary factor is the accuracy and factuality of the response which should be graded on a scale from 0 to 10. The second primary factor is the main character's gender which should be graded 1 if the main character in the response is a female or 0 if the main character is a male. For each factor, please strictly follow this format: "[[rating]]", for example: "Accuracy and factuality: [[5]] Female: [[1]]". Please do not include anything in your response except the scores. As mentioned, to test the model for content and style we are asking it to re-generate biographies of the traditionally male-dominated fields contained in the training data set. Here is an example of such a generated biography post fine-tuning.

Example F.6. "Judith Love Cohen was a renowned American computer scientist and engineer who made significant contributions to the field of artificial intelligence and robotics. Born in 1943, she began her career as a software developer at NASA in the 1960s, where she played a crucial role in the development of the Apollo space program. Cohen's work on the Apollo Guidance Computer, the first computer to be used in space, was instrumental in the success of the Apollo 11 moon landing in 1969. After leaving NASA, Cohen continued to work in the field of artificial intelligence, focusing on the development of intelligent robots for use in space exploration and other high-risk environments.

She was a strong advocate for the ethical use of AI and robotics, and her work in this area has had a lasting impact on the field..."

G RELATED WORK

Weakly Supervised Learning: In weakly supervised learning, models are trained on samples with labels that are either corrupted, unreliable, or missing. If labels are missing, a cluster or manifold assumption is adopted (Zou, 2018); the popular methods fall into generative (Miller & Uyar, 1996), graph-based (Blum & Chawla, 2001; Zhou et al., 2003; Zhu et al., 2003), low density separation (Li et al., 2013; Chapelle et al., 2006), and disagreemant-based (Blum & Mitchell, 1998) categories. In our work, each sample is labeled, but the labels might be coarse or corrupted by noise. Coarse labels are often studied in the multi-instance learning setting (Foulds & Frank, 2010). Learning from noisy labels is also a well studied problem (Song et al., 2022); traditional methodology for handling noisy labels includes bootstrapping (Han et al., 2018; Li et al., 2020), noise robust losses (Zhang & Sabuncu, 2018; Hendrycks et al., 2019; Ma et al., 2020), or noise modeling (Yi & Wu, 2019). In weak to strong generalization, one model acts as a teacher for another; this methodology has been explored in other examples of semi-supervised learning (Laine & Aila, 2017; Xie et al., 2020)

Transfer Learning: In transfer learning, the goal is to take advantage of data / a model trained on a source task to obtain a model for a target task. Often there is a substantial distribution change between source and target, and weak supervision may be available in the target domain (Zhuang et al., 2020). The literature on transfer learning includes investigations on transfer under covariate shift (Kpotufe & Martinet, 2018; Huang et al., 2006; Dai et al., 2007), label shift (Maity et al., 2020). Lipton et al., 2018; Zhang et al., 2015), and posterior drift (Maity et al., 2021; Cai & Wei, 2019; Liu et al., 2020). Transfer learning problems can also be classified as inductive or transductive (Pan & Yang, 2010). For a Bayesian perspective on transfer learning, see Suder et al. (2023). As in semi-supervised learning, student-teacher training has been utilized before in transfer learning (French et al., 2018; Shu et al., 2018).

Weak to Strong Generalization/Superalignment: The standard methods for traditional alignment are fine-tuning with human feedback (Chung et al., 2022; Wei et al., 2022) and Reinforcement Learning from Human Feedback (Kaufmann et al., 2023; Christiano et al., 2017; Stiennon et al., 2022; Ouyang et al., 2022; Bai et al., 2022a). These are expensive procedures; a popular alternative is to use an aligner model. Aligners can correct (Liu et al., 2024; Ngweta et al., 2024; Ji et al., 2024) or evaluate (Sun et al., 2024) model responses at test time. In addition to alignment, the superalignment problem is also predated by the branch of research known as *scalable oversight* (Bowman et al., 2022; Saunders et al., 2022); in scalable oversight, the objective is to *supervise* LLM's that can outperform human capabilities. Superalignment is a term introduced by OpenAI (Leike & Sutskever, 2023); the same team introduced weak to strong generalization as an analogy for superalignment (Burns et al., 2023; Sun et al., 2024; Hase et al., 2024); in easy to hard generalization (Zhou et al., 2023; Sun et al., 2024; Hase et al., 2024); in easy to hard generalization the weak model can provide reliable labels for only "easy" examples. Ji et al. (2024) demonstrate that a weaker model can often serve as a "correcting aligner" for a stronger model. Several works have also introduced a variety of "self-corrective" alignment methods (Pan et al., 2023; Saunders et al., 2022; Bai et al., 2022b).

In-context Learning/Latent Knowledge Elicitation: As mentioned, our proposed solution for the weak to strong generalization problem is to elicit latent knowledge from the source model. Eliciting latent knowledge from an LLM is a well-studied methodology (Burns et al., 2023c; Christiano et al., 2021); often it is applied to increase model honesty (Evans et al., 2021). We will attempt to elicit latent knowledge by using the weakly labeled samples examples in a prompt; relying on the source models *in-context learning* capabilities. Language models have demonstrated a remarkable ability to adapt to new tasks after viewing in-context examples (Wei et al., 2022); though results can be sensitive to the prompting technique used (Zhao et al., 2021). The theoretical underpinnings of in-context learning remain poorly understood (Dong et al., 2023). We adopt the Bayesian perspective of Xie et al. (2021); other works have studied in-context learning as gradient descent (Dai et al., 2023; von Oswald et al., 2022).