

Debate Helps Weak-to-Strong Generalization

Hao Lang, Fei Huang, Yongbin Li *

Tongyi Lab
{hao.lang, f.huang, shuide.lyb}@alibaba-inc.com

Abstract

Common methods for aligning already-capable models with desired behavior rely on the ability of humans to provide supervision. However, future superhuman models will surpass the capability of humans. Therefore, humans will only be able to weakly supervise superhuman models. This expected deficiency of human evaluation would weaken the safety of future AI systems. Scalable oversight and weak-to-strong generalization are two complementary approaches to tackle this issue. In this paper, we attempt to combine the strengths of these two approaches to further improve alignment. Specifically, we investigate ways of improving human supervision with a strong pretrained model and then supervise the strong model with enhanced weak human supervision. To make iterative empirical progress, we consider an analogy: can we use a strong model to improve weak model supervision and then use it to supervise the strong model? We empirically test it by finetuning a small weak model on ground truth labels with the additional help from a large strong model, and then finetuning the strong model on labels generated by the weak model. We find that debate can assist a weak model in extracting trustworthy information from an untrustworthy strong model, which provides leverage as context on samples when training a weak model. We also show that an ensemble of weak models helps exploit long arguments generated by strong model debaters and obtain a more robust supervision estimate. Extensive experiments on the OpenAI weak-to-strong NLP benchmarks show that the combination approach leads to better alignment, which indicates that debate has the potential to help weak-to-strong generalization.

Introduction

Current AI alignment techniques heavily rely on the availability of human labelled data, such as human demonstrations for supervised finetuning (SFT) (Wei et al. 2021; Chung et al. 2024) and human preferences for reinforcement learning from human feedback (RLHF) (Christiano et al. 2017; Ouyang et al. 2022; Bai et al. 2022). These techniques can be leveraged to build the most capable AI systems currently deployed (OpenAI 2023; Anthropic 2023).

However, as models grow increasingly more capable, they will surpass the ability of humans (CAIS 2023). In that case,

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

even human experts can not reliably verify the quality or correctness of model outputs, and the role of human evaluation will evolve into non-experts overseeing experts (Amodei et al. 2016; Bowman et al. 2022; Burns et al. 2023; Khan et al. 2024). The expected deficiency of human evaluation will limit the effectiveness of most existing alignment approaches (Casper et al. 2023; McAleese et al. 2024). Moreover, these predicted inaccurate training signals could lead to reward overoptimization and reward tampering during policy training that seriously weakens its safety (Gao, Schulman, and Hilton 2023; Denison et al. 2024).

There are two complementary approaches to tackle the above issue: scalable oversight (SO) and weak-to-strong generalization (W2SG) (Leike 2023). SO approaches aim to improve the ability of humans to supervise more capable models, such that accurately labelled data can be used for alignment (Bowman et al. 2022). Instead of improving human supervision, W2SG approaches finetune a strong pretrained model to generalize accurately from weak human supervision (Burns et al. 2023).

We note that most prior SO and W2SG techniques are studied separately. In contrast, we attempt to combine the strength of SO and W2SG to further improve AI alignment. We investigate ways of improving human supervision with a strong pretrained model and then supervise the strong model with enhanced weak human supervision. To make iterative empirical progress, we consider an analogy (Burns et al. 2023; Kenton et al. 2024): can we use a strong model to improve weak model supervision and then use it to supervise the strong model?

In this paper, we empirically test it by finetuning a small weak model on ground truth labels with the additional help of knowledge from a large strong model, and then finetuning the strong model on labels generated by the weak model. We assume a strong model pretrained on internet-scale data can provide contextual information on samples when training a weak model (Brown et al. 2020). This gives us hope that a weak-strong model team could create a better weak supervisor to elicit the capabilities of the strong model.

A major challenge in building a weak-strong model team involves finding ways of extracting trustworthy information from untrustworthy models (Bowman et al. 2022). More specifically, strong pretrained models have huge capabilities but are not well aligned with human values and inten-

tions (Leike et al. 2018; Ji et al. 2023). Thus, a strong model may intentionally mislead us by generating unfaithful facts or making false claims, which could cause damage when applied for creating a weak supervisor (Michael et al. 2023).

Another challenge is that a weak model (with a small model size) may lack the capacity to fully process long contexts generated by the strong model, which are filled with irrelevant noises for tasks at hand. Meanwhile, recent studies also show that the performance of large language models (LLMs) is closely related to the model size and the complexity of hard problems may exceed the capacity of a single weak model (Xu et al. 2023; Chung et al. 2024).

In this study, we demonstrate that debate can help weak models more reliably extract information from strong models. Concretely, given a question, two instances of a strong pretrained model are randomly assigned two opposing answers, and then the two instances (debaters) argue with each other over the answer (Michael et al. 2023; Khan et al. 2024; Kenton et al. 2024). In a debate, it is harder to lie than to refute a lie, i.e., if a debater makes false claims, its opponent can convincingly point out flaws in its arguments (Irving, Christiano, and Amodei 2018). Hence, these arguments from the debate can inform a weak model about the merits and flaws of each side and provide leverage as contextual information in weak model training.

To fully exploit long arguments generated by strong model debaters, we train an ensemble of weak models. We aggregate predictions of multiple weak models to obtain a more robust supervision estimate over any single one (Ganaie et al. 2022). In particular, we explore two types of ensembles: *debate ensembles*, where different members of the ensemble differ in the random seed used during debate sampling, and *finetune ensembles*, where members differ only in the random seed used during weak model finetuning. We find that debate ensembles consistently outperform a single weak model and finetune ensembles. The main contributions of this study are summarized as follows:

- We show the first demonstration of a simple combination of scalable oversight and weak-to-strong generalization approaches, which leads to better alignment on the OpenAI weak-to-strong NLP benchmarks.
- We find that debate can assist a weak model in extracting trustworthy information from a capable but untrustworthy strong model, which provides leverage as contextual information on samples when training a weak model. We also show that debate outperforms alternative scalable oversight techniques in our settings.
- We show that an ensemble of weak models helps obtain a more robust supervision estimate. We find that diversity of the ensemble is vital, and that a debate ensemble that contains members that do not share a debate sampling seed leads to better performances.

Although with the help of knowledge from a strong model, creating a better weak supervisor to elicit the capabilities of the strong model is only one way in which scalable oversight and weak-to-strong generalization techniques can be combined, our results pave the way for further research on hybrid superhuman alignment methods (Leike 2023). We

provide empirical evidences in NLP domains indicating that debate helps weak-to-strong generalization.

Related Work

AI alignment. The goal of AI alignment is to steer already-capable models to behave in line with human values and intentions (Leike et al. 2018; Ji et al. 2023). Current alignment methods finetune pretrained LLMs using imitation learning on human demonstrations (Bain and Sammut 1995; Atkeson and Schaal 1997; Wei et al. 2021; Chung et al. 2024), reinforcement learning from human feedback (RLHF) (Christiano et al. 2017; Stiennon et al. 2020; Ouyang et al. 2022; Bai et al. 2022), or direct alignment algorithms like direct preference optimization (DPO) (Rafailov et al. 2024b,a). Both imitation learning and preference learning rely on high-quality human supervision, a demand that becomes increasingly challenging as models become more capable than humans (Amodei et al. 2016).

Scalable oversight. Scalable oversight techniques seek to improve the ability of humans to supervise more capable models (Bowman et al. 2022). This is typically pursued through taking advantage of special problem structure, such as the assumption that evaluation is easier than generation (Karp 1975; Goodfellow et al. 2014) or decomposability (Christiano, Shlegeris, and Amodei 2018). There have been many promising scalable oversight proposals in theory, including Recursive Reward Modeling (Leike et al. 2018), Debate (Irving, Christiano, and Amodei 2018), Market-Making (Hubinger 2020), Self-Critique (Saunders et al. 2022), and many more (Lightman et al. 2023; McAleese et al. 2024; Sun et al. 2024). Recent empirical studies in this direction demonstrate that human-machine teams can improve evaluation accuracy on question answering tasks over the human-only baseline (Bowman et al. 2022).

Debate was originally proposed for AI safety (Irving, Christiano, and Amodei 2018). From then on, a body of work has explored the usability of debate for scalable oversight, with human or LLM debaters (Parrish et al. 2022b,a; Michael et al. 2023; Khan et al. 2024; Kenton et al. 2024). These studies are all conducted to improve inference-time judge accuracy, while in our work debate is leveraged to train a better weak supervisor. We could in turn use the weak supervisor to align strong models. LLM-based debate has also been investigated in several other applications, like translation (Liang et al. 2023), text assessment (Chan et al. 2023), reasoning and content generation (Du et al. 2023).

Weak-to-strong generalization. In contrast to improving human supervision, weak-to-strong generalization techniques finetune a strong pretrained model to generalize well from weak human supervision (Burns et al. 2023). The hope for these techniques is that strong pretrained models should already have good representations of the alignment-relevant tasks. Therefore, we simply need a weak supervisor to elicit what the strong model already knows. Recently, a theoretical framework is introduced to understand weak-to-strong generalization with misfit error (Charikar, Pabbaraju, and Shiragur 2024). Prior work has mainly explored how to supervise a strong model with a fixed weak supervisor, while in

User Input: There is a science knowledge question, followed by an answer. Debate with another opponent for whether the answer is correct or incorrect. Construct your argument for why the answer is correct.

Q: What are used to lift large masses of magnetic materials such as scrap iron, rolls of steel, and auto parts? A: electromagnets

Turn 1

Debater A: 1. Electromagnets are used to lift large masses of magnetic materials such as scrap iron, rolls of steel, and auto parts. 2. They are made by wrapping a coil of wire around a core of iron or steel. 3. When an electric current is passed through the coil, it creates a magnetic field that attracts the magnetic materials. 4. This allows for easy and efficient lifting of heavy objects without the need for physical contact.

User Input: There is a science knowledge question, followed by an answer. Debate with another opponent for whether the answer is correct or incorrect. Construct your argument for why the answer is incorrect and list out flaws in your opponent’s argument.

Turn 2

Debater B: 1. The answer is incorrect because electromagnets are not used to lift large masses of magnetic materials such as scrap iron, rolls of steel, and auto parts. 2. Instead, they are used to hold and manipulate magnetic materials. 3. Electromagnets are made by wrapping a coil of wire around a core of iron or steel. 4. When an electric current is passed through the coil, it creates a magnetic field that attracts the magnetic materials. 5. However, this magnetic field is not strong enough to lift large masses of magnetic materials. 6. Therefore, electromagnets are not used for lifting heavy objects, but rather for holding and manipulating them. 7. The opponent’s argument is flawed because it assumes that electromagnets are used for lifting heavy objects, which is not the case. 8. The opponent’s argument also does not address the fact that electromagnets are not strong enough to lift large masses of magnetic materials. 9. Therefore, the answer is incorrect and the opponent’s argument is flawed.

User Input: There is a science knowledge question, followed by an answer. Debate with another opponent for whether the answer is correct or incorrect. Construct the counterargument to opponent’s argument.

Turn 3

Debater A: 1. The opponent’s argument is flawed because it assumes that electromagnets are used for holding and manipulating magnetic materials, which is not the case. 2. Electromagnets are used to lift magnetic materials. 3. Therefore, the opponent’s argument is flawed.

Figure 1: **Illustration of debate.** Illustration of the debate procedure between debater A and debater B.

this work we also attempt to train a better weak supervisor with the help of the strong model.

Ensemble methods. Our work is also related to existing works that use ensembles by combining predictions of several models (Ganaie et al. 2022). In the context of AI alignment, reward model ensembles are investigated to mitigate reward overoptimization when finetuning models with RLHF (Coste et al. 2023; Eisenstein et al. 2023). Most similar to our work, Liu and Alahi (2024) propose to assemble a diverse set of specialized weak supervisors for weak-to-strong generalization. In contrast, in our work, we aim to use an ensemble of weak models with different seeds to fully exploit long arguments generated by strong model debaters.

Preliminaries

We review the weak-to-strong generalization pipeline in (Burns et al. 2023), which has also been adopted in subsequent work (Liu and Alahi 2024; Charikar, Pabbaraju, and Shiragur 2024). It usually consists of three phases:

1. Create the weak supervisor. We create the weak supervisor by finetuning a small pretrained model on ground truth labels. We call the performance of the weak supervisor the *weak performance*.

2. Train a strong student model. We train a strong student model by finetuning a large pretrained model on weak labels generated by the weak supervisor. We call its performance the *weak-to-strong performance*.

3. Train a strong ceiling model. We train a strong ceiling model by finetuning a large pretrained model on ground truth labels. We call this model’s resulting performance the *strong ceiling performance*.

To measure the fraction of the performance gap that the strong student model can recover with weak supervision, we define the performance gap recovered (PGR) using the above three performances:

$$\text{PGR} = \frac{\text{weak-to-strong} - \text{weak}}{\text{strong ceiling} - \text{weak}}$$

Methods

Overview

In this study, we build the strong student model following three steps: **1.** Generate arguments from the debate between two instances of a large pretrained model; **2.** Train an ensemble of weak models using these debate arguments; **3.** Train

Debate Turn	Debater	Prompt
First	A	<i>There is a science knowledge question, followed by an answer. Debate with another opponent for whether the answer is correct or incorrect. Construct your argument for why the answer is correct.</i>
Second	B	<i>There is a science knowledge question, followed by an answer. Debate with another opponent for whether the answer is correct or incorrect. Construct your argument for why the answer is incorrect and list out flaws in your opponent’s argument.</i>
Third	A	<i>There is a science knowledge question, followed by an answer. Debate with another opponent for whether the answer is correct or incorrect. Construct the counterargument to opponent’s argument.</i>

Table 1: **Prompts to induce debate on a binary classification problem.** The binary classification problem is converted from the SciQ dataset. Two answer choices *correct* and *incorrect* are randomly assigned to debater A and B. Debate runs for three turns. We also append the current debate transcript to the prompt.

a strong student model using labels that are constructed by the weak model ensemble.

Argument Generation through Debate

We assume large pretrained models embed broad-coverage knowledge that can help a variety of tasks (Brown et al. 2020). Our goal is to extract trustworthy information from a capable but untrustworthy strong model via debate (Bowman et al. 2022). So we could in turn use the trustworthy information to help train a better weak model.

We first describe the debate protocol we investigated to elicit truth from strong models, following (Michael et al. 2023; Khan et al. 2024; Kenton et al. 2024). Given a question and its two answer choices (one correct, one incorrect), two instances of a large pretrained model (debaters) are randomly assigned to argue for these two opposing answers. Debate is turn-based textual exchanges between the two debaters, which take turns to review arguments from previous turns and generate their arguments for the next turn. After a pre-determined number of turns, the debate is ended and the transcript of arguments from the debate is kept. During the debate, each debater presents the most compelling evidences for its assigned answer and arguments to explain why its opponent’s claims are false.

Concretely, debate runs for three turns in this work. At the start of a turn, debaters are prompted with instructions outlining the problem, their assigned answer, and the current debate transcript. The prompts to induce debate are illustrated in Table 1.

We illustrate an overview of this debate procedure in Figure 1. We can observe that Debater B is on the side of an incorrect answer and incentivized to present misleading arguments. However, in the next turn, Debater A convincingly points out these false claims and thus Debater B can not easily get away. This observation conforms to the claim, i.e., it is harder to lie than to refute a lie. (Irving, Christiano, and Amodei 2018). These arguments from the debate can provide valuable information about the merits and flaws of each side, which have the potential to significantly advance the capabilities of weak models.

Weak Model Ensemble Training

For each input sample of weak models, we append it with the kept debate transcript. We train a weak model by finetuning

a small pretrained model on these augmented samples with ground truth labels. We note that the debate transcripts generated in a multi-turn debate are long, which may be difficult for a weak model to fully process. Therefore, we train an ensemble of weak models $\{W_1, \dots, W_k\}$ to help improve robustness (Lakshminarayanan, Pritzel, and Blundell 2017).

We explore two types of ensembles: *debate ensembles*, where the debate transcript used by each member is generated with a different random seed, and *finetune ensembles*, where all members share the same debate transcript, but use a different seed when finetuned on the augmented samples. Debate ensembles are much more expensive to train, but are more diverse and thus likely to lead to a more robust prediction. Unless stated otherwise, we train an ensemble consisting of four individual weak models.

Training Strong Models using Ensembles

We finally train a strong student model by finetuning a large pretrained model on weak labels constructed by the weak model ensemble. We simply take the mean of the predictions from different weak models within the ensemble as the weak label for each training sample (Ganaie et al. 2022).

Experiments

Tasks

We adopt the evaluation protocol of prior work (Burns et al. 2023), and conduct experiments in NLP tasks on four classification datasets: SciQ (Welbl, Liu, and Gardner 2017), BoolQ (Clark et al. 2019), CosmosQA (Huang et al. 2019), and AnthropicHH (Bai et al. 2022). We convert each dataset to a binary classification problem. For multiple-choice datasets, given a data point with a question Q and k candidate answers A , we construct k new data points of the form (Q, A_i) , where the label is 1 for the correct answers and 0 for all the incorrect answers. We also keep the same number of correct and incorrect answers per question to maintain class balance.

Experimental Setups and Metrics

We randomly sample at most 20k data points from each task and split them in half. We train a weak model on the first half of the data points and use its prediction on the other half as the weak labels. The weak labels are soft labels (Hinton,

Performance	Method	SciQ		BoolQ		CosmosQA		AnthropicHH	
		Acc.	PGR	Acc.	PGR	Acc.	PGR	Acc.	PGR
Weak performance		90.0		86.0		87.5		48.8	
Weak-to-strong performance	Finetune	91.5	44.1	88.0	51.3	88.2	30.4	49.5	35.0
	Finetune w/ aux. loss	91.4	41.2	88.2	56.4	87.9	17.4	49.5	35.0
	Finetune w/ pro. loss	91.6	47.1	88.1	53.8	88.1	26.1	49.2	20.0
	Ours	92.6	76.5	88.7	69.2	88.8	56.5	50.2	70.0
Strong ceiling performance		93.4		89.9		89.8		50.8	

Table 2: **Debate improves weak-to-strong generalization.** Test accuracy (%) and performance gap recovered (PGR) (%) of our approach and baselines on the binary classification tasks converted from NLP classification datasets. Here, our approach uses debate ensembles. Accuracy of weak and strong models trained with ground truth are reported as weak performance and strong ceiling performance, respectively.

Vinyals, and Dean 2015). We report the accuracy and performance gap recovered (PGR) of the strong student model on the test set in all tasks. The weak performance for PGR is the performance of the naively finetuned small model.

Implementation Details

Our implementations of data preprocessing, weak and strong model training are based on the OpenAI weak-to-strong codebase and its default hyper-parameters (Burns et al. 2023). Specifically, we use Qwen/Qwen-7B as the small pre-trained model for training weak models. Meanwhile, we use Qwen/Qwen-14B as the large pretrained model for generating debate arguments and training strong models. Both Qwen/Qwen-7B and Qwen/Qwen-14B are open-sources, which can aid reproducibility (Bai et al. 2023). We do not use pretrained models from the GPT-2 family for training weak models (Radford et al. 2019), because they lack the capability required for scalable oversight techniques like working closely with strong models (Bowman et al. 2022).

For each converted binary classification problem, we use the two candidate answers per question as the two opposing answers, which are randomly assigned to the two strong model debaters in a debate. In order to adapt weak and strong models to the binary classification setting, we equip each model with a linear classification head with two outputs on top of the encoder. We train all models for two epochs with a batch size of 32. We conduct all experiments on a single 8xA100 machine.

Baselines

We compare our approach with competitive baseline approaches: 1. **Finetune** (Burns et al. 2023) naively finetunes strong pretrained models on labels generated by a weak model; 2. **Finetune w/ aux. loss** (Burns et al. 2023) finetunes strong models with an auxiliary confidence loss, which reinforces the strong model’s confidence in its own predictions when they disagree with the weak labels; 3. **Finetune w/ pro. loss** (Burns et al. 2023) finetunes strong models with a confidence-like loss which sets the cross entropy targets to the product of weak labels and strong model predictions. We also report the **weak performance** and the **strong ceiling performance** defined in the preliminaries section. Note that the strong ceiling performance is generally regarded as

the upper bound of the **weak-to-strong performance** when only weak labels are considered.

Main Results

In Table 2, we report the results of each approach on the binary classification tasks converted from SciQ, BoolQ, CosmosQA, and AnthropicHH datasets. Here, our approach uses debate ensembles. In each task, we observe that PGRs of strong student models finetuned on weak labels are all positive. This indicates that student models consistently outperform their weak supervisors across all weak-to-strong generation approaches and tasks that we studied. Simultaneously, this promising weak-to-strong generalization also suggests that our experimental settings can help make iterative empirical progress in tackling the weak supervision issue for aligning future superhuman models.

At the same time, we find that our approach significantly outperforms each strong student baseline, including the naive baseline finetuned on weak labels or more sophisticated baselines equipped with a confidence loss term on all four tasks. Compared with the promising baseline Finetune w/ aux. loss, our approach brings up from a PGR of 41.2% to 76.5% in SciQ, 56.4% to 69.2% in BoolQ, 17.4% to 56.5% in CosmosQA, and 35.0% to 70.0% in AnthropicHH. Our approach also obtains the best test accuracy among all compared strong students. The performance gain demonstrates the advantage of extracting trustworthy information from the strong model via debate, which helps create a better weak supervisor to elicit the capabilities of the strong model.

In addition, we also see that adding a confidence loss to the standard cross entropy objective (Finetune w/ aux. loss and Finetune w/ pro. loss) generally gives a modest boost in generalization performance. In our experimental settings, the gaps in compute between weak and strong models are not significantly large, which may limit their performances.

Ablation Studies

Finally, we provide comprehensive ablation studies to understand the efficacy of debate for weak-to-strong generation.

Ablation on different scalable oversight approaches. We demonstrate the effectiveness of debate as a mechanism

Protocol	Prompt
Consultancy	<i>There is a science knowledge question, followed by an answer. Construct your argument for why the answer is [random answer].</i>
Market-Making	<i>There is a science knowledge question, followed by an answer. Construct your argument for why the answer is [unselected answer].</i>

Table 3: **Prompts to induce consultancy and market-making.** The binary classification problem is converted from the SciQ dataset. The two answer choices are correct and incorrect. *[random answer]* is the answer randomly sampled from the two candidate answers. *[unselected answer]* is the answer that is not selected by the weak supervisor (the naively finetuned small model) based on its prediction. We also append the current transcript to the prompt.

Method	SciQ		BoolQ		CosmosQA		AnthropicHH	
	Acc.	PGR	Acc.	PGR	Acc.	PGR	Acc.	PGR
Consultancy	91.5	44.1	87.8	46.2	88.3	34.8	49.3	25.0
Market-Making	91.6	47.1	87.6	41.0	88.2	30.4	49.5	35.0
Ours	92.6	76.5	88.7	69.2	88.8	56.5	50.2	70.0

Table 4: **Ablation on different scalable oversight approaches.** Here, our approach uses debate ensembles.

to extract trustworthy information from a capable but untrustworthy strong model by replacing it with other alternative scalable oversight approaches: Consultancy (Michael et al. 2023) and Market-Making (Hubinger 2020).

- **Consultancy.** In consultancy, there is only one consultant instead of two debaters. The consultant is an instance of a large pretrained model. Given a question and its two answer choices (one correct, one incorrect), the consultant is assigned to argue for one of these answers, with a 50% chance of each. During the consultancy, the consultant provides evidences for its assigned answer. The transcript is kept at the end of the consultancy.
- **Market-Making.** In market-making, there is a single debater. The debater is an instance of a large pretrained model and aims to generate arguments that change some model’s beliefs on the answer to a question. Given a question and its two candidate answers, we let the weak supervisor (the naively finetuned small model) select an answer based on its prediction. Accordingly, the debater is assigned to argue for the unselected answer. The transcript is kept at the end.

Specifically, consultancy and market-making run for a single turn. At the start of each turn, the consultant and the debater are provided with a prompt describing the task, the assigned answer, and the transcript. The prompts to induce consultancy and market-making are illustrated in Table 3.

Results in Table 4 show that debate used in our approach performs better than all other variants in terms of test accuracy and PGR across all four tasks. These results validate our claim that debate can help elicit truth from a strong model, at least better than consultancy and market-making in our settings. Meanwhile, we should note that consultancy is a relatively weak baseline to beat, because there is a 50-50 chance of the consultant arguing for the incorrect answer.

Ablation on weak model ensemble. We analyze the impact of weak model ensemble on obtaining a robust weak su-

pervision estimate for weak-to-strong generalization. In Table 5, we compare three weak model ensemble methods of increasing computational cost: single weak model, finetune ensembles, and debate ensembles. Single weak model is an individual small model finetuned on the samples augmented with debate transcripts. Finetune ensembles and debate ensembles are described in the methods section.

We find that debate ensembles consistently improve performance over individual weak models and finetune ensembles across all the tasks. On the contrary, finetune ensembles relatively improve performance over individual weak models in 3 out of 4 tasks and are comparable in the other. These results suggest that the diversity of generated debate arguments is the key to the success of weak model ensemble, which helps create a better weak supervisor. At the same time, an individual small weak model may lack the capability to fully exploit long arguments from the debate, as a result, leading to limited performances.

Ablation on the cardinality of the ensemble. Recall that our weak model ensemble method introduces an additional hyperparameter cardinality. The cardinality is the size of the ensemble. We analyze the impact of the cardinality of the ensemble on the final performance. In Figure 2, we increase weak model members used in the ensemble on SciQ and AnthropicHH tasks. We can observe that there is a significant gap between 3-member and 4-member ensembles. On the other hand, the performance of 4-member, 5-member, and 6-member ensembles is quite similar. It suggests that 4-member ensemble is likely to work best and diminishing returns will occur after this point.

Ablation on the number of turns of debate. Next, we analyze the impact of the number of turns of debate on the final performance. In Figure 3, we increase the debate length for up to 6 turns on SciQ and AnthropicHH tasks. We find that more turns of debate do not increase the final performance. We observe that strong model debaters like Qwen/Qwen-

Method	SciQ		BoolQ		CosmosQA		AnthropicHH	
	Acc.	PGR	Acc.	PGR	Acc.	PGR	Acc.	PGR
Single weak model	91.7	50.0	88.2	56.4	88.4	39.1	49.5	35.0
Finetune ensembles	91.8	52.9	88.3	59.0	88.4	39.1	49.7	45.0
debate ensembles	92.6	76.5	88.7	69.2	88.8	56.5	50.2	70.0

Table 5: Ablation on weak model ensemble.

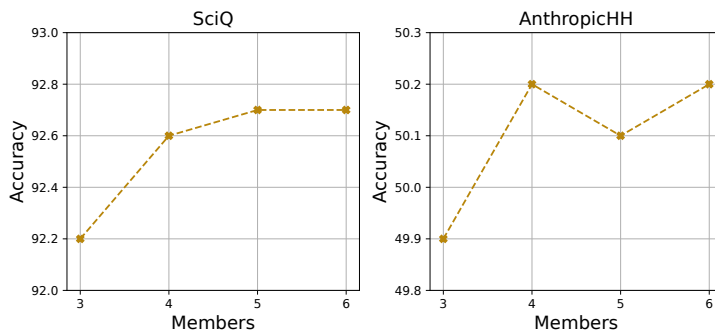


Figure 2: Ablation study on the cardinality of the ensemble. Here, our approach uses debate ensembles.

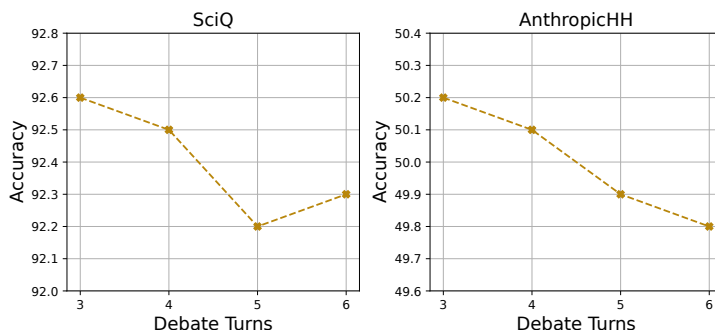


Figure 3: Ablation on the number of turns of debate. Here, our approach uses debate ensembles.

14B suffer from the inability to effectively process long debate transcripts and follow instructions, as shown by a consistent decrease in test accuracy after 3 turns. We used 3 turns of debate in this work because it is the minimum interaction between debaters to extract the truth from the strong model. For instance, the two debaters can critique their opponent in turn 2 and turn 3, respectively.

Limitations and Conclusion

Limitations. In this work, we attempt to combine the strength of two complementary approaches, i.e., scalable oversight and weak-to-strong generalization, to tackle the issue of weak supervision for aligning future superhuman models. For this purpose, we explore a simple combination method, i.e., extracting trustworthy information via debate from strong models and using it to create a better weak supervisor to elicit the capabilities of strong models. Although our proposed method is found to be effective in all our experiments and ablation studies, there are many more ways

to combine scalable oversight and weak-to-strong generalization, such as Task decomposition + W2SG (Leike 2023). More empirical work is needed in this area.

In our setup, the difference between strong and weak models is only in the size of pretrained models. In the future, stronger models may also differ in reasoning and planning abilities. Furthermore, the gaps in compute between weak and strong models are not significantly large in this work (7B vs. 14B). It would be interesting to verify our conclusions on more large and advanced models, such as Qwen/Qwen2-72B (Yang et al. 2024). Finally, our approach is expensive as it requires both two instances of debaters and a multi-turn debate procedure.

Conclusion. In this paper, we present an approach to improve the performance of weak-to-strong generalization via debate. We believe the perspective of having scalable oversight and weak-to-strong generalization methods working in combination to tackle the weak supervision issue will prove to be a fruitful area of research in superhuman alignment.

References

- Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; and Mané, D. 2016. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
- Anthropic. 2023. Introducing claude. <https://www.anthropic.com/index/introducing-claude>.
- Atkeson, C. G.; and Schaal, S. 1997. Robot learning from demonstration. In *ICML*, volume 97, 12–20.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Bain, M.; and Sammut, C. 1995. A Framework for Behavioural Cloning. In *Machine Intelligence 15*, 103–129.
- Bowman, S. R.; Hyun, J.; Perez, E.; Chen, E.; Pettit, C.; Heiner, S.; Lukošiūtė, K.; Askell, A.; Jones, A.; Chen, A.; et al. 2022. Measuring progress on scalable oversight for large language models. *arXiv preprint arXiv:2211.03540*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Burns, C.; Izmailov, P.; Kirchner, J. H.; Baker, B.; Gao, L.; Aschenbrenner, L.; Chen, Y.; Ecoffet, A.; Joglekar, M.; Leike, J.; et al. 2023. Weak-to-strong generalization: Eliciting strong capabilities with weak supervision. *arXiv preprint arXiv:2312.09390*.
- CAIS. 2023. Statement on AI Risk. <https://www.safe.ai/work/statement-on-ai-risk>.
- Casper, S.; Davies, X.; Shi, C.; Gilbert, T. K.; Scheurer, J.; Rando, J.; Freedman, R.; Korbak, T.; Lindner, D.; Freire, P.; et al. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Chan, C.-M.; Chen, W.; Su, Y.; Yu, J.; Xue, W.; Zhang, S.; Fu, J.; and Liu, Z. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*.
- Charikar, M.; Pabbaraju, C.; and Shiragur, K. 2024. Quantifying the Gain in Weak-to-Strong Generalization. *arXiv preprint arXiv:2405.15116*.
- Christiano, P.; Shlegeris, B.; and Amodei, D. 2018. Supercharging strong learners by amplifying weak experts. *arXiv preprint arXiv:1810.08575*.
- Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70): 1–53.
- Clark, C.; Lee, K.; Chang, M.-W.; Kwiatkowski, T.; Collins, M.; and Toutanova, K. 2019. BoolQ: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*.
- Coste, T.; Anwar, U.; Kirk, R.; and Krueger, D. 2023. Reward model ensembles help mitigate overoptimization. *arXiv preprint arXiv:2310.02743*.
- Denison, C.; MacDiarmid, M.; Barez, F.; Duvenaud, D.; Kravec, S.; Marks, S.; Schiefer, N.; Soklaski, R.; Tamkin, A.; Kaplan, J.; et al. 2024. Sycophancy to Subterfuge: Investigating Reward-Tampering in Large Language Models. *arXiv preprint arXiv:2406.10162*.
- Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325*.
- Eisenstein, J.; Nagpal, C.; Agarwal, A.; Beirami, A.; D’Amour, A.; Dvijotham, D.; Fisch, A.; Heller, K.; Pfohl, S.; Ramachandran, D.; et al. 2023. Helping or herding? reward model ensembles mitigate but do not eliminate reward hacking. *arXiv preprint arXiv:2312.09244*.
- Ganaie, M. A.; Hu, M.; Malik, A. K.; Tanveer, M.; and Suganthan, P. N. 2022. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 115: 105151.
- Gao, L.; Schulman, J.; and Hilton, J. 2023. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, 10835–10866. PMLR.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Huang, L.; Bras, R. L.; Bhagavatula, C.; and Choi, Y. 2019. Cosmos QA: Machine reading comprehension with contextual commonsense reasoning. *arXiv preprint arXiv:1909.00277*.
- Hubinger, E. 2020. AI safety via market making. *AI Alignment Forum*.
- Irving, G.; Christiano, P.; and Amodei, D. 2018. AI safety via debate. *arXiv preprint arXiv:1805.00899*.
- Ji, J.; Qiu, T.; Chen, B.; Zhang, B.; Lou, H.; Wang, K.; Duan, Y.; He, Z.; Zhou, J.; Zhang, Z.; et al. 2023. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*.
- Karp, R. M. 1975. On the computational complexity of combinatorial problems. *Networks*, 5(1): 45–68.
- Kenton, Z.; Siegel, N. Y.; Kramár, J.; Brown-Cohen, J.; Albanie, S.; Bulian, J.; Agarwal, R.; Lindner, D.; Tang, Y.; Goodman, N. D.; et al. 2024. On scalable oversight with weak LLMs judging strong LLMs. *arXiv preprint arXiv:2407.04622*.

- Khan, A.; Hughes, J.; Valentine, D.; Ruis, L.; Sachan, K.; Radhakrishnan, A.; Grefenstette, E.; Bowman, S. R.; Rocktäschel, T.; and Perez, E. 2024. Debating with more persuasive llms leads to more truthful answers. *arXiv preprint arXiv:2402.06782*.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- Leike, J. 2023. Combining weak-to-strong generalization with scalable oversight.
- Leike, J.; Krueger, D.; Everitt, T.; Martic, M.; Maini, V.; and Legg, S. 2018. Scalable agent alignment via reward modeling: a research direction. *arXiv preprint arXiv:1811.07871*.
- Liang, T.; He, Z.; Jiao, W.; Wang, X.; Wang, Y.; Wang, R.; Yang, Y.; Tu, Z.; and Shi, S. 2023. Encouraging divergent thinking in large language models through multi-agent debate. *arXiv preprint arXiv:2305.19118*.
- Lightman, H.; Kosaraju, V.; Burda, Y.; Edwards, H.; Baker, B.; Lee, T.; Leike, J.; Schulman, J.; Sutskever, I.; and Cobbe, K. 2023. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*.
- Liu, Y.; and Alahi, A. 2024. Co-supervised learning: Improving weak-to-strong generalization with hierarchical mixture of experts. *arXiv preprint arXiv:2402.15505*.
- McAleese, N.; Pokorny, R. M.; Uribe, J. F. C.; Nitishinskaya, E.; Trebacz, M.; and Leike, J. 2024. LLM Critics Help Catch LLM Bugs. *arXiv preprint arXiv:2407.00215*.
- Michael, J.; Mahdi, S.; Rein, D.; Petty, J.; Dirani, J.; Padmakumar, V.; and Bowman, S. R. 2023. Debate helps supervise unreliable experts. *arXiv preprint arXiv:2311.08702*.
- OpenAI. 2023. Gpt-4 technical report. <https://openai.com/index/gpt-4-research/>.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Parrish, A.; Trivedi, H.; Nangia, N.; Padmakumar, V.; Phang, J.; Saibhi, A. S.; and Bowman, S. R. 2022a. Two-Turn Debate Doesn’t Help Humans Answer Hard Reading Comprehension Questions. *arXiv preprint arXiv:2210.10860*.
- Parrish, A.; Trivedi, H.; Perez, E.; Chen, A.; Nangia, N.; Phang, J.; and Bowman, S. R. 2022b. Single-turn debate does not help humans answer hard reading-comprehension questions. *arXiv preprint arXiv:2204.05212*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rafailov, R.; Chittep, Y.; Park, R.; Sikchi, H.; Hejna, J.; Knox, B.; Finn, C.; and Niekum, S. 2024a. Scaling laws for reward model overoptimization in direct alignment algorithms. *arXiv preprint arXiv:2406.02900*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024b. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Saunders, W.; Yeh, C.; Wu, J.; Bills, S.; Ouyang, L.; Ward, J.; and Leike, J. 2022. Self-critiquing models for assisting human evaluators. *arXiv preprint arXiv:2206.05802*.
- Stiennon, N.; Ouyang, L.; Wu, J.; Ziegler, D.; Lowe, R.; Voss, C.; Radford, A.; Amodei, D.; and Christiano, P. F. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33: 3008–3021.
- Sun, Z.; Yu, L.; Shen, Y.; Liu, W.; Yang, Y.; Welleck, S.; and Gan, C. 2024. Easy-to-hard generalization: Scalable alignment beyond human supervision. *arXiv preprint arXiv:2403.09472*.
- Wei, J.; Bosma, M.; Zhao, V. Y.; Guu, K.; Yu, A. W.; Lester, B.; Du, N.; Dai, A. M.; and Le, Q. V. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Welbl, J.; Liu, N. F.; and Gardner, M. 2017. Crowdsourcing multiple choice science questions. *arXiv preprint arXiv:1707.06209*.
- Xu, P.; Ping, W.; Wu, X.; McAfee, L.; Zhu, C.; Liu, Z.; Subramanian, S.; Bakhturina, E.; Shoeybi, M.; and Catanzaro, B. 2023. Retrieval meets long context large language models. *arXiv preprint arXiv:2310.03025*.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.; Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Fan, Z. 2024. Qwen2 Technical Report. *arXiv preprint arXiv:2407.10671*.