# SCALABLE OVERSIGHT VIA PARTITIONED HUMAN SUPERVISION

**Anonymous authors**Paper under double-blind review

000

001

002003004

010 011

012

013

014

016

018

019

021

024

025

026

027

028

029

031

033 034

037

038

040

041

042

043

046

047

048

051

052

#### **ABSTRACT**

As artificial intelligence (AI) systems approach and surpass expert human performance across a broad range of tasks, obtaining high-quality human supervision for evaluation and training becomes increasingly challenging. Our focus is on tasks that require deep knowledge and skills of multiple domains. Unfortunately, even the best human experts are knowledgeable only in a single narrow area, and will not be able to evaluate the correctness of advanced AI systems on such superhuman tasks. However, based on their narrow expertise, humans may provide a weak signal, i.e., a complementary label indicating an option that is incorrect. For example, a cardiologist could state that "this is not related to cardiology," even if they cannot identify the true disease. Based on this weak signal, we propose a scalable oversight framework that enables us to evaluate frontier AI systems without the need to prepare the ground truth. We derive an *unbiased* estimator of top-1 accuracy from complementary labels and quantify how many complementary labels are needed to match the variance of ordinary labels. We further introduce two estimators to combine scarce ordinary labels with abundant complementary labels. We provide finite-sample deviation guarantees for both complementary-only and the mixed estimators. Empirically, we show that we can evaluate the output of large language models without the ground truth, if we have complementary labels. We further show that we can train an AI system with such weak signals: we show how we can design an agentic AI system automatically that can perform better with this partitioned human supervision.

#### 1 Introduction

As foundation models and artificial intelligence (AI) systems (OpenAI, 2025a; Anthropic, 2025; DeepSeek, 2025; DeepMind, 2025) approach and in some areas surpass expert human performance, supervision itself becomes a key bottleneck. Current alignment pipelines such as supervised fine-tuning (SFT), reinforcement learning from human feedback (RLHF; Ziegler et al. (2019); Stiennon et al. (2020); Ouyang et al. (2022)), or reinforcement learning from verifiable rewards (RLVR; Shao et al. (2024); Yu et al. (2025); Lambert et al. (2025)) presuppose that humans can reliably provide supervision or design verifiers for training and evaluation. Yet for the superalignment regime (Bowman et al., 2022), future models will tackle problems whose solutions are too technical or too cross-disciplinary for any single human to verify comprehensively. When we cannot produce ground truth or prepare automated verifiers, how should we *evaluate* and *train* such systems?

We begin from an observation about the expertise in high-skill domains in the human society (Wuchty et al., 2007). As tasks grow in difficulty, human experts specialize ever more narrowly: cardiology vs. oncology in medicine, sector-specific analysts in finance, or sub-subfields in science. Specialists can often make *local* determinations with high precision, e.g., a cardiologist might say, "this is *not* a cardiology case," while an oncologist might point out, "this pattern is never seen in the field of oncology." While such judgments may fail to positively certify correctness, they readily certify that certain options are wrong. In other words, even when positive ground truth is out of reach, human experts can often provide reliable *complementary labels* (Ishida et al., 2017; Yu et al., 2018; Wang et al., 2024a; 2025): labels indicating one or more options that are definitely the incorrect.

This paper develops a scalable oversight (Amodei et al., 2016) framework that treats such *partitioned human supervision* as a first-class supervision primitive. Our starting point is a multi-choice



Figure 1: Superhuman tasks are routed to a human expert who can only give a weak signal that they are not capable of solving the task. We show that this weak signal, i.e., *complementary label*, can be used to evaluate/train an AI system.

evaluation setting where ordinary (true) labels may be scarce or unavailable, but complementary labels arise naturally from partitioned expertise and are abundant. We show that one can utilize these complementary labels to propose an unbiased estimator of top-1 accuracy for any system under test, and combine them with scarce ordinary labels for refinement of the estimator. The resulting estimators allow us to (i) *evaluate* frontier models without preparing any ground truth, and (ii) use these evaluations as learning signals to *design* agentic systems on tasks beyond expert human abilities.

Our work complements many other active lines of research on scalable oversight. For example, weak-to-strong generalization (Burns et al., 2024) studies whether strong models trained on weak (noisy) supervisors can exceed those supervisors, possibly by eliciting knowledge that the strong model already possesses. Easy-to-hard generalization (Sun et al., 2024) trains evaluators on easier problems and applies them to supervise harder ones, by leveraging the assumption that evaluation is easier than generation. Debate (Khan et al., 2024) elicits truth via adversarial arguments that a judge selects among. Reinforcement learning from AI feedback or Constitutional AI (Bai et al., 2022) replaces human preferences with AI-generated feedback. Recursive task decomposition (Wu et al., 2021) breaks the problem into leaf subtasks that humans can directly label or compare, for tasks that are easily decomposable. We address a new, orthogonal axis: when human experts cannot supply correct answers for hard, cross-domain instances, can we use their partitioned "this option is wrong" judgments to obtain performance estimates and useful training signals?

We formalize a labeling protocol in which each instance is routed to a randomly chosen domain specialist responsible for exactly one option (class). The specialist either confirms that option (yielding an ordinary label) or confidently rejects it (yielding a complementary label). Under this uniform wrong-index design, we derive an *unbiased* estimator of accuracy from complementary labels alone via a simple linear correction, and characterize its variance. We then introduce two *mixture estimators* that combine ordinary and complementary labels: (i) an inverse-variance—weighted (IVW) estimator with a practical plug-in form, and (ii) a closed-form maximum-likelihood estimator (ML).

We empirically validate the framework in three settings. First, we perform *statistical validation* on many popular large language model (LLM) benchmarks, which confirms that we do not need the ground truth to evaluate the performance of AI systems experimentally, such as LLMs. We focus on *real-world evaluation* on tasks that cannot be solved by a single expert human with only narrow expertise. As a proof-of-concept, we use finance and medical benchmarks to demonstrate that partitioned feedback from sector analysts or highly-specialized doctors enable accurate model evaluation without the ability to solve the task alone. Finally, we show that this weak signal can be used as a training signal for an AI system. More specifically, we use complementary labels for *agentic training*: we replace ordinary accuracy with our estimator as the fitness signal inside agent search pipelines, and show that this improves downstream performance, demonstrating a pathway to training when only complementary feedback is available.

Our main contributions are as follows: 1) We introduce scalable oversight via partitioned human supervision, a protocol that exploits real-world specialization to collect *complementary labels* at scale for superhuman tasks. 2) We derive an unbiased complementary-label estimator for top-1 accuracy, analyze its variance, and show how many complementary labels are needed to match the variance of the estimator obtained from ordinary labels. We further propose two mixture estimators together with finite-sample bounds and practical plug-in implementations. 3) We demonstrate empirically that these estimators enable both evaluation and training without the full ground truth.

### 2 METHODOLOGY

In this section, we explain our problem setup and notations and explain our proposed method.

**Setup and Notations** We consider a multiple-choice evaluation task with K options per item (e.g., K=4 for choices A, B, C, and D). Formally, each item is represented by an input  $x\in\mathcal{X}$ , where  $\mathcal{X}$  denotes the input domain (e.g., the space of question texts). Each item has a correct label  $y\in\{1,\ldots,K\}$ , representing the index of the correct option after any shuffling. We call this the ordinary label. Let (x,y) denote a random pair drawn from an identical joint distribution  $\mathcal{D}$  over  $\mathcal{X}\times\{1,\ldots,K\}$ .

Equivalently, we can regard to first sample item x, and then a human expert provides label y based on the underlying conditional distribution p(y|x). An AI system (e.g., a machine learning classifier, a single LLM, or an LLM agent/workflow) produces a prediction  $\hat{y} = f(x) \in \{1, \dots, K\}$ . For an item with an ordinary label y, define the indicator  $Z = \mathbb{I}\{\hat{y} = y\} \in \{0, 1\}$ . The (population) top-1 accuracy is then  $A = \Pr(Z = 1) = \Pr(f(x) = y)$ , where  $\Pr(\cdot)$  denotes probability with respect to the joint distribution  $\mathcal{D}$ .

We assume the complementary label is drawn uniformly at random among the non-y classes (Ishida et al., 2017):

$$p(\bar{y} = k \mid y) = \frac{1}{K - 1}, \quad \forall k \neq y. \tag{1}$$

That is, for each item, the revealed wrong index is sampled uniformly from the K-1 non-y classes. For an item with a complementary label  $\bar{y} \neq y$ , define  $W = \mathbb{I}\{\hat{y} \neq \bar{y}\} \in \{0,1\}$ . The assumption in Eq. 1 is enforced by the following data collection protocol. We assume that we have a set of questions  $\{x_i\}_{i=1}^N$  and K annotators, where each annotator is a human expert in the field of one of the classes. For each question  $x_i$ , choose an annotator randomly (let's say an annotator in charge of the k-th class is chosen), and ask "Is the answer of question  $x_i$  class k?". If the answer is yes, we save this in the ordinary-label set as an item with (ordinary) class k label. If the answer is no, we save this in the complementary-label set as an item with (complementary) class k label.

Finally, we will have two disjoint evaluation sets, drawn from the same distribution and evaluated by the same system: 1) An *ordinary-label set* of size  $n_{\rm o}$ , yielding indicators  $Z_1,\ldots,Z_{n_{\rm o}}$  and total correct count  $S_{\rm o}=\sum_{j=1}^{n_{\rm o}}Z_j$ . 2) complementary-label set of size  $n_{\rm c}$ , yielding indicators  $W_1,\ldots,W_{n_{\rm c}}$  and total "weakly-correct" count  $S_{\rm c}=\sum_{i=1}^{n_{\rm c}}W_i$ . The total number of samples is  $N=n_{\rm o}+n_{\rm c}$ .

#### 2.1 ESTIMATOR FROM COMPLEMENTARY LABELS

It is straightforward to estimate the system accuracy A by using ordinary labels as  $\widehat{A}_{\mathrm{ord}} = \frac{S_o}{n_o}$ . If we want to use complementary labels to estimate A, a naive approach would be to count the proportion of "avoiding the complementary label," i.e.,  $\widehat{q} = \frac{1}{n_c} \sum_{i=1}^{n_c} W_i$ . However,  $\widehat{q}$  is not an unbiased estimator of the accuracy. Thus, we propose the following complementary-label estimator:

**Corollary 1.** *Under the assumption in Eq. (1), the estimator* 

$$\hat{A}_{\text{comp}} = (K-1)\hat{q} - (K-2)$$
 (2)

is unbiased for A, where  $K \geq 3$  is the number of choices.

We show the proof in Appendix A. This linear correction is precisely the 0–1 loss specialization of the risk-rewrite identity of Eqs. (8) and (10) in Ishida et al. (2019); replacing expectations by sample means yields our unbiased estimator. A full derivation is provided in Appendix B.

Even if we have an unbiased estimator with only complementary labels, the variance is expected to become larger compared with the case when we use the same number of ordinary labels. This is because complementary labels are "weaker" than ordinary labels in the sense that K-1 different complementary labels is equivalent to 1 ordinary label. We first confirm this as follows. Since W is Bernoulli with mean  $q = \frac{A+K-2}{K-1}$ , we obtain

$$\operatorname{Var}(\widehat{A}_{\text{comp}}) = (K-1)^2 \operatorname{Var}\left(\frac{1}{n_c} \sum_{i=1}^{n_c} W_i\right) = \frac{(K-1)^2}{n_c} q(1-q) = \frac{(A+K-2)(1-A)}{n_c}.$$
(3)

<sup>&</sup>lt;sup>1</sup>Here "weakly-correct" means that the model prediction is *not* equal to the complementary label. For example, when K=4, avoiding the complementary label still leaves a  $\frac{2}{3}$  chance of being wrong.

This expression depends on the unknown A. In practice, a plug-in estimator is used. For the ordinary-label dataset, the number of correct predictions satisfies  $S_o \sim \text{Bin}(n_o, A)$ , where Bin(n, p) denotes the binomial distribution with parameters n (number of independent trials) and p (success probability in each trial). Thus,  $\mathbb{E}[S_o/n_o] = A$ , and  $\text{Var}(S_o/n_o) = A(1-A)/n_o$ , where  $\mathbb{E}[\cdot]$  and  $\text{Var}(\cdot)$  denote expectation and variance, respectively, taken with respect to the underlying population distribution  $\mathcal{D}$  over (x,y). Comparing this with Eq. (3), the variance ratio is

$$\frac{\operatorname{Var}(\widehat{A}_{\text{comp}})}{\operatorname{Var}(S_{\text{o}}/n_{\text{o}})} = \frac{(A+K-2)\,n_{\text{o}}}{A\,n_{\text{c}}}.\tag{4}$$

Hence, when  $n_{\rm c}=n_{\rm o}$ , the complementary estimator has larger variance for all  $K\geq 3$ , as expected. Next, it would be interesting to derive how much more complementary labels we will need to prepare to have the same amount of variance as the ordinary label case. To match the variance of the ordinary estimator, the size of the complementary dataset must be

$$n_{\rm c} = \left(1 + \frac{K - 2}{A}\right) n_{\rm o}.\tag{5}$$

That is, the larger A is, the fewer additional complementary labels are needed. Since A is unknown, one can use a small ordinary-labeled set to bound A and infer the required  $n_{\rm c}$ , motivating the combined estimator in § 2.2.

#### 2.2 COMBINING ORDINARY AND COMPLEMENTARY LABELS

We now combine information from both datasets. Recall  $\widehat{A}_{\mathrm{ord}} = \frac{S_o}{n_o}$  and  $\widehat{A}_{\mathrm{comp}} = (K-1)\frac{S_c}{n_c} - (K-2)$  are both unbiased under the independent and identically distributed (i.i.d.) sampling assumption.

Inverse Variance Weighted Estimator (IVW) Consider the linear combination  $\widehat{A}_{\text{mix}} = w \, \widehat{A}_{\text{ord}} + (1-w) \, \widehat{A}_{\text{comp}}$  where  $w \in [0,1]$ . Because the two components are independent and unbiased (disjoint samples),  $\widehat{A}_{\text{mix}}$  is unbiased with variance

$$\operatorname{Var}(\widehat{A}_{\text{mix}}) = w^2 \frac{A(1-A)}{n_0} + (1-w)^2 \frac{(K-1)^2}{n_0} q(1-q), \quad q = \frac{A+K-2}{K-1}.$$

Minimizing it w.r.t. w yields

$$w^* = \frac{\frac{(K-1)^2}{n_c} q(1-q)}{\frac{A(1-A)}{n_o} + \frac{(K-1)^2}{n_c} q(1-q)}.$$
 (6)

The corresponding minimum variance equals  $\operatorname{Var}_{\min}(\widehat{A}_{\min}) = \frac{\frac{A(1-A)}{n_o} \cdot \frac{(K-1)^2}{n_c} q(1-q)}{\frac{A(1-A)}{n_o} + \frac{(K-1)^2}{n_c} q(1-q)}$ . Since A is

unknown, we use plug-in estimators:  $\widehat{\mathrm{Var}}(\widehat{A}_{\mathrm{ord}}) = \frac{\widehat{A}_{\mathrm{ord}}(1-\widehat{A}_{\mathrm{ord}})}{n_{\mathrm{o}}}$ , and  $\widehat{\mathrm{Var}}(\widehat{A}_{\mathrm{comp}}) = \frac{(K-1)^2}{n_{\mathrm{c}}}\,\widehat{q}(1-\widehat{q})$ ,  $\widehat{q} = \frac{S_{\mathrm{c}}}{n_{\mathrm{c}}}$ . Then

$$\widehat{w} = \frac{\widehat{\text{Var}}(\widehat{A}_{\text{comp}})}{\widehat{\text{Var}}(\widehat{A}_{\text{ord}}) + \widehat{\text{Var}}(\widehat{A}_{\text{comp}})}, \text{ where } \widehat{A}_{\text{IVW}} = \widehat{w} \, \widehat{A}_{\text{ord}} + (1 - \widehat{w}) \, \widehat{A}_{\text{comp}}.$$
(7)

Its (plug-in) variance is the harmonic mean:

 $\widehat{\mathrm{Var}}(\widehat{A}_{\mathrm{IVW}}) = \widehat{\mathrm{Var}}(\widehat{A}_{\mathrm{ord}}) \, \widehat{\mathrm{Var}}(\widehat{A}_{\mathrm{comp}}) / (\widehat{\mathrm{Var}}(\widehat{A}_{\mathrm{ord}}) + \widehat{\mathrm{Var}}(\widehat{A}_{\mathrm{comp}})). \text{ Substituting } q = \frac{A + K - 2}{K - 1} \text{ into Eq. (6) gives}$ 

$$w^* = \frac{(A+K-2)\,n_0}{A\,n_c + (A+K-2)\,n_0},\tag{8}$$

and replacing A by a pilot  $\tilde{A}$  (e.g.  $\tilde{A} = \hat{A}_{ord}$ ), yields a closed-form estimator of  $w^*$ . If we prepare a separate set other than  $n_c + n_o$  samples to estimate the weight, we can maintain the unbiasedness of  $\hat{A}_{mix}$ ; if not (as in our experiments), we will show in § 2.3 that we have a consistent estimator.

Closed-Form Maximum-Likelihood Estimator (ML) The ordinary set induces a likelihood term in A given by  $S_{\rm o} \sim {\rm Bin}(n_{\rm o},A)$ . The complementary set induces a likelihood term  $S_{\rm c} \sim {\rm Bin}(n_{\rm c},q)$  with  $q = \frac{A+K-2}{K-1}$ . Assuming independence between the two datasets, the joint log-likelihood is

$$\ell(A) = S_0 \log A + (n_0 - S_0) \log(1 - A) + S_c \log q + (n_c - S_c) \log(1 - q).$$

Setting  $\partial \ell/\partial A=0$  yields the score equation  $\frac{S_{\rm o}}{A}-\frac{n_{\rm o}-S_{\rm o}}{1-A}+\frac{S_{\rm c}}{A+K-2}-\frac{n_{\rm c}-S_{\rm c}}{1-A}=0$ , which simplifies to a quadratic form in A. Let  $T_{\rm o}=n_{\rm o}-S_{\rm o}$ ,  $T_{\rm c}=n_{\rm c}-S_{\rm c}$ ,  $\alpha=N$ ,  $\beta=(K-2)(T_{\rm o}+T_{\rm c})+(K-3)S_{\rm o}-S_{\rm c}$ ,  $\gamma=-(K-2)S_{\rm o}$ . Then the ML is the unique root in [0,1]:

$$\widehat{A}_{\rm ML} = \frac{-\beta + \sqrt{\beta^2 - 4\alpha\gamma}}{2\alpha}.$$
 (9)

A large-sample standard error is  $\operatorname{se}(\widehat{A}_{\mathrm{ML}}) \approx \left[\frac{n_{\mathrm{o}}}{\widehat{A}_{\mathrm{ML}}(1-\widehat{A}_{\mathrm{ML}})} + \frac{n_{\mathrm{c}}}{(K-1)^2\,\widehat{q}(1-\widehat{q})}\right]^{-1/2}$ , based on observed or expected information.

If  $n_{\rm c}=0$ , Eq. (9) reduces to  $\widehat{A}_{\rm ord}$ . If  $n_{\rm o}=0$ , it reduces to  $\widehat{A}_{\rm comp}$ . For K=4, one always has  $\widehat{q}\in \left[\frac{2}{3},1\right]$  when  $A\in [0,1]$ , which bounds the complementary-component variance. Moreover, a one-step Newton update around the truth shows that the IVW estimator with optimal fixed weights coincides with the ML estimator up to the first order<sup>2</sup>, and both are statistically efficient.

#### 2.3 THEORETICAL ANALYSIS

We analyze the deviation of the complementary-label estimator  $\widehat{A}_{\text{comp}}$  under the uniform wrongindex assumption Eq. (1). Although  $\widehat{A}_{\text{comp}}$  is unbiased, its variance can be non-negligible at small sample sizes. We present two finite-sample deviation bounds: a distribution-free Hoeffding bound (Boucheron et al., 2003) and a data-dependent Bernstein bound (Maurer & Pontil, 2009) leveraging empirical variance. We then extend these to a fixed-weight mixture estimator.

**Bounds for**  $\widehat{A}_{\text{comp}}$  We derive a unified bound for  $\widehat{A}_{\text{comp}}$  by taking the minimum of the Hoeffding and empirical Bernstein inequalities. This formulation is both variance-free (via Hoeffding) and variance-adaptive (via Bernstein). The proof is shown in Appendix E.

**Theorem 2.** With probability at least  $1 - \delta$ ,

$$\left| \widehat{A}_{\text{comp}} - A \right| \; \leq \; (K-1) \, \min \left\{ \sqrt{\frac{\log(2/\delta)}{2 \, n_{\text{c}}}}, \; \sqrt{\frac{2 \, \widehat{q}(1-\widehat{q})}{n_{\text{c}}-1} \, \log \frac{4}{\delta}} + \frac{7 \, \log(4/\delta)}{3 \, (n_{\text{c}}-1)} \right\}.$$

The first branch (Hoeffding) is simple and variance-free but conservative. The second branch (empirical Bernstein) adapts to the empirical variance  $\hat{q}(1-\hat{q})$  and is typically tighter when  $\hat{q}$  is close to 0 or 1. Both bounds require only i.i.d. sampling of  $\{W_i\}$  under Eq. (1), and hold for any  $n_c$ .

Bounds for  $\hat{A}_{\text{mix}}$  Consider the mixture estimator with a (possibly data-dependent) weight  $w \in [0,1]$ :  $\hat{A}_{\text{mix}} = w \, \hat{A}_{\text{ord}} + (1-w) \, \hat{A}_{\text{comp}}$  where  $\hat{A}_{\text{ord}} = \frac{1}{n_{\text{o}}} \sum_{j=1}^{n_{\text{o}}} Z_j$ , and  $Z_j \in \{0,1\}$ . Since the ordinary and complementary sets are disjoint, we have  $\hat{A}_{\text{mix}} - A = w(\hat{A}_{\text{ord}} - A) + (1-w)(\hat{A}_{\text{comp}} - A)$ . With a union bound and Hoeffding inequality, we can derive the following proposition.

**Proposition 3.** For any split  $\delta = \delta_0 + \delta_c$ , with  $\delta_0, \delta_c > 0$ , with probability at least  $1 - \delta$ ,

$$|\widehat{A}_{\text{mix}} - A| \leq \min \left\{ w \sqrt{\frac{\log(2/\delta_{\text{o}})}{2 n_{\text{o}}}} + (1 - w)(K - 1) \sqrt{\frac{\log(2/\delta_{\text{c}})}{2 n_{\text{c}}}}, \\ w \left( \sqrt{\frac{2 \widehat{A}_{\text{ord}} (1 - \widehat{A}_{\text{ord}})}{n_{\text{o}} - 1} \log \frac{4}{\delta_{\text{o}}}} + \frac{7 \log(4/\delta_{\text{o}})}{3 (n_{\text{o}} - 1)} \right) + (1 - w) \left( \sqrt{\frac{2 (K - 1)^2 \widehat{q} (1 - \widehat{q})}{n_{\text{c}} - 1} \log \frac{4}{\delta_{\text{c}}}} + \frac{7 (K - 1) \log(4/\delta_{\text{c}})}{3 (n_{\text{c}} - 1)} \right) \right\}.$$

A symmetric choice  $\delta_{\rm o} = \delta_{\rm c} = \delta/2$  yields the simpler logs  $\log(4/\delta)$  and  $\log(8/\delta)$  respectively.

<sup>&</sup>lt;sup>2</sup>A detailed proof is provided in Appendix D.

The union-bound approach provides guarantees that hold for *any* weight  $w \in [0,1]$ , including data-dependent plug-in choices such as  $\widehat{w}_{\mathrm{IVW}}$ . This makes it directly applicable in practice, although the constants are typically looser than those obtained from the moment generating function based Bernstein bound (See Theorem 4).

We now state a Bernstein-type PAC bound for the mixed estimator that *complements* the union-bound inequalities established in Proposition 3; its proof is deferred to Appendix F.

**Theorem 4.** With probability at least  $1 - \delta$ , the following holds for  $\widehat{A}_{mix}$ :

$$\left| \widehat{A}_{\text{mix}} - A \right| \le \sqrt{2v \log \frac{2}{\delta}} + c \log \frac{2}{\delta}, \tag{10}$$

where 
$$v = w^2 \frac{A(1-A)}{n_o} + (1-w)^2 \frac{(K-1)^2 q(1-q)}{n_c}$$
, and  $c = \max\left\{\frac{w}{n_o}, \frac{(1-w)(K-1)}{n_c}\right\}$ .

This bound tightens Hoeffding's inequality by incorporating both the variance term v and the (range-driven) linear term c. It is valid for any fixed choice of the weight  $w \in [0,1]$ , for instance if w is set as a constant or estimated on an independent pilot set. If instead w is estimated from the same evaluation data (e.g. the plug-in  $\widehat{w}_{\text{IVW}}$ ), the fixed-weight assumption is violated; in that case one must either revert to the union-bound bounds (Proposition 3), or apply standard remedies such as sample splitting (Chernozhukov et al., 2018; Wager & Athey, 2018)or a finite grid with a union bound (Boucheron et al., 2013; Shalev-Shwartz & Ben-David, 2014). For practical plug-in forms and further discussion, see Appendix G.

#### 3 EXPERIMENTS

We evaluate our unbiased and mixed estimators in three settings: (I) Statistical validation on standard multiple-choice benchmarks, where we verify unbiasedness and variance properties (see Eqs. (2), (4), (16)); (II) Practical evaluation on two real-world classification benchmarks: the Japanese financial dataset EDINET-Bench (Sugiura et al., 2025), where each option corresponds to a professional industry domain and annotators provide weak complementary feedback; and the English Medical Abstracts dataset (Schopf et al., 2023), which covers multiple disease categories with natural clinical text. (III) Agentic training in agent-based workflows, where we replace ordinary accuracy with our proposed estimator as the fitness function and examine the resulting agent performance. For agent scaffolding, we use ADAS (Hu et al., 2025) and AFlow (Zhang et al., 2025c). For details of the labeling protocol, please refer to Appendix C.

Statistical Validation of Complementary and Mixed Estimators Statistical validation is conducted on four benchmarks. Since oracle labels are available, complementary labels are synthetically constructed following Appendix C. We report three types of evaluation metrics. (A) Unbiasedness and variance are examined by subsampling 300 ordinary and 300 complementary labels (120 for GPQA due to dataset size). Accuracy is estimated using Eq. (2), and variability is quantified through standard deviations, reported as  $\widehat{A} \pm \widehat{\sigma}$  with  $\widehat{\sigma}$  denoting the estimated standard deviation. (B) The relative efficiency of complementary labels is assessed by computing the number of complementary samples required to achieve the same variance as  $n_0$  ordinary labels, according to Eq. (4), where the full ordinary dataset is used to approximate the ground-truth accuracy A. (C) To mimic practical annotation protocols (Appendix C), we assume that each ordinary label can be replaced by (K-1) complementary labels, where K is the number of answer options. This setting is used to evaluate the mixed estimators in Eq. (7) and Eq. (9).

Each dataset is evaluated on all five metrics using gpt-5-nano (OpenAI, 2025b), with three independent runs averaged; aggregated outcomes are summarized in Table 1. For the MATH-MC benchmark, we use GPT-4.1-nano instead of GPT-5-nano. This choice is due to (i) output-format constraints in GPT-5-nano that conflict with our final-answer extraction protocol, and (ii) the need to avoid ceiling effects: in our pilot runs, GPT-5-nano attains near-ceiling accuracy on MATH-MC, which compresses the observable gaps among estimators.

Table 1 summarizes the results across benchmarks, together with a reference evaluation using all available ordinary labels (Ord-Eval). Importantly, higher accuracy in isolation does not imply a better estimator; the desideratum is to approximate the Ord-Eval reference as closely and reliably

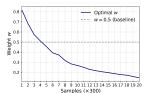
Table 1: Performance of estimators across benchmarks. Values are reported as mean accuracy  $\pm$  standard deviation across three random seeds; the average of per-run standard deviations is additionally shown in parentheses. Detailed estimator setups and confidence-interval analyses are provided in Appendix I. The last column reports the macro-average across datasets (counting Math and Math-COT separately). For each benchmark, the estimator whose mean accuracy is closest to the Ord-Eval reference is highlighted in bold.

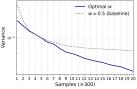
Estimator	MMLU-Pro	MedQA-USMLE	GPQA	$\mathbf{MATH}^{\dagger}$	MATH(CoT) <sup>‡</sup>	Average
Ord	$78.33 \pm 1.73  (\pm 2.38)$	$92.89 \pm 1.35 (\pm 1.48)$	$64.17 \pm 1.67 (\pm 4.39)$	$47.56 \pm 3.91 (\pm 2.88)$	$84.89 \pm 0.77 (\pm 2.07)$	73.57
Comp-n <sub>o</sub>	$77.00 \pm 12.49 (\pm 7.95)$	$92.67 \pm 1.53 (\pm 2.67)$	59.17 $\pm$ 3.82 ( $\pm$ 9.42)	$48.44 \pm 10.78  (\pm 7.72)$	$80.44 \pm 2.78 (\pm 4.98)$	71.54
Comp-Var	$75.67 \pm 2.15 (\pm 2.51)$	$90.61 \pm 1.43 (\pm 1.69)$	$63.67 \pm 5.01 (\pm 4.28)$	$41.10 \pm 3.17 (\pm 2.93)$	$81.35 \pm 0.29 (\pm 2.28)$	70.48
IVW-0.5	$77.89 \pm 1.58 (\pm 1.80)$	$91.61 \pm 1.11 (\pm 1.15)$	$65.28 \pm 1.34 (\pm 3.33)$	$43.44 \pm 3.95 (\pm 2.52)$	$83.56 \pm 1.17 (\pm 1.58)$	72.36
IVW	77.97 $\pm$ 1.58 ( $\pm$ 1.79)	$91.86 \pm 1.11 (\pm 1.13)$	$65.14 \pm 1.38 \ (\pm 3.30)$	$44.87 \pm 3.82 (\pm 2.37)$	83.86 $\pm$ 0.83 ( $\pm$ 1.56)	72.74
ML	$77.94 \pm 1.58 (\pm 1.79)$	$91.65 \pm 1.08 (\pm 1.18)$	$65.11 \pm 1.38 \ (\pm 3.28)$	44.75 $\pm$ 3.79 ( $\pm$ 2.36)	$83.65 \pm 1.04 (\pm 1.59)$	72.62
Ord-Eval	77.97	92.66	59.52	44.21	83.89	-

as possible under scarce supervision. With finite ordinary labels and (K-1)-times complementary labels (where K is the number of choices), Ord and Comp- $n_o$  can occasionally yield point estimates that are numerically closer to Ord-Eval (as seen in MedQA and GPQA). However, their standard errors are substantially wider, indicating poor reliability and limited reproducibility. For Comp-Var, increasing the number of complementary labels reduces variance, consistent with Eq. (3) and Eq. (4). We report both across-seed variability and the average within-run standard deviation, the latter more directly reflecting estimator variance. Ord-Eval itself should not be regarded as the ground-truth accuracy—especially on GPQA, where dataset size is small and uncertainty remains substantial—but rather as an empirical reference.

Among mixture estimators, IVW consistently improves over the equal-weighted baseline in withinrun standard deviation (IVW-0.5), confirming that inverse-variance weighting provides the optimal linear combination of ordinary and complementary labels. ML further achieves narrower intervals owing to its joint likelihood formulation. However, ML requires more precise pilot estimates of accuracy and may thus be less robust in practice, whereas IVW remains stable by directly incorporating variance estimates. This distinction is supported by the within-run standard deviations, which indicate that IVW achieves competitive efficiency with simpler computation.

Overall, while Comp- $n_o$  may occasionally appear closest to Ord-Eval in mean accuracy, its instability makes it unreliable for reproducible evaluation. IVW and ML instead strike the best balance between bias and variance, yielding stable and reproducible estimates close to the Ord-Eval reference across datasets. The empirical closeness of the two methods further validates our theoretical analysis, suggesting that mixture estimators can serve as a practical foundation for scalable oversight in settings where ground-truth labels are scarce or infeasible.





(a) Optimal w vs. number of samples .

(b) Corresponding variance on a log scale.

Figure 2: Comparison of optimal weighting and baseline (w=0.5). Optimal weighting consistently achieves lower variance compared to the equal-weighted baseline, demonstrating more reliable estimation with limited samples.

Ablation Study on IVW. Recall that the

optimal coefficient for IVW is given in Eq. (8). When setting  $n_{\rm c}=(K-1)n_{\rm o}$  (where K denotes the number of choices), the expression simplifies to  $\frac{A+K-2}{AK+K-2}=\frac{1}{K}\left(1+\frac{(K-1)(K-2)}{AK+K-2}\right)$ . Since  $0 \le A \le 1$ , the coefficient lies in the range  $\frac{1}{2} \le \frac{1}{K}\left(1+\frac{(K-1)(K-2)}{AK+K-2}\right) \le 1$ . In particular, when A is close to 1—as suggested by the Ord-Eval reference—the optimal weight approaches  $\frac{1}{2}$ , explaining why in Table 1 the IVW estimator often performs similarly to the ordinary baseline.

To further examine whether the closed-form coefficient yields the optimal trade-off, we conduct an ablation on **MedQA-USMLE** by varying the number of complementary labels from  $n_{\rm c}=n_{\rm o}$  to  $n_{\rm c}=20n_{\rm o}$  (with  $n_{\rm o}=300$ ). Results are averaged over three random seeds. As shown in Figure 2, when  $n_{\rm c}\approx 3n_{\rm o}$ —the setting used in Table 1—the optimal weight is close to the fixed baseline w=0.5. However, for  $n_{\rm c}< n_{\rm o}$  or  $n_{\rm c}>6n_{\rm o}$ , the optimal variance is consistently lower than the baseline. This behavior is intuitive: once the number of complementary labels exceeds the variance-

Table 2: Mean acc.  $\pm$  standard deviation across 3 seeds, w/ the average of per-run standard deviations. Bold: Lowest variance across 3 runs. See App. I for details.

Estimator	EDINET	EDINET (Extended)	Medical Abstract
Ord	$19.35 \pm 5.59 (\pm 7.13)$	$21.67 \pm 2.46 (\pm 2.89)$	$70.38 \pm 0.62 \ (\pm 0.96)$
Comp-(K-1)	$10.75 \pm 8.12 (\pm 16.46)$	$26.60 \pm 3.45 (\pm 5.86)$	$68.79 \pm 0.71 (\pm 1.13)$
IVW-0.5	$15.05 \pm 5.18 (\pm 8.98)$	$24.14 \pm 2.26 \ (\pm 3.27)$	$69.58 \pm 0.57 (\pm 0.74)$
IVW	$17.95 \pm 4.88  (\pm 6.53)$	$22.64 \pm 2.27 (\pm 2.59)$	<b>69.71</b> $\pm$ <b>0.56</b> ( $\pm$ 0.73)
ML	17.93 $\pm$ 4.87 ( $\pm$ 6.29)	22.61 $\pm$ 2.28 ( $\pm$ 2.26)	$69.70 \pm 0.56 (\pm 0.74)$

matching threshold in Eq. (4), the complementary estimator alone has smaller variance, and the closed-form IVW solution naturally shifts weight away from ordinary labels. By contrast, the fixed baseline cannot adapt, leading to consistently higher variance.

Benchmark Evaluation in Real-World Tasks To demonstrate the practicality of our method in real-world scenarios, we conduct proof-of-concept experiments for *estimation and learning under partitioned human supervision*, as directly collecting superhuman tasks would be prohibitively costly. Following the data collection protocol in Algorithm 1, for each domain we randomly select an expert and ask whether a given object belongs to their field. Applied to the benchmarks, complementary labels are generated from the oracle labels by splitting each instance into 1 ordinary label and K-1 complementary labels, corresponding to the K candidate classes.

We evaluate on two benchmarks. First, we use EDINET-Bench (Sugiura et al., 2025), a Japanese financial report classification dataset, where each industry corresponds to a professional domain, and we adopt the prompting setup of Sugiura et al. (2025). Since the original dataset contains only 496 samples (comparable to GPQA), we extend it following their released code, resulting in 3,269 samples; we refer to this as EDINET-Bench Extended. To our knowledge, there are no public medical benchmarks in which each class explicitly corresponds to a medical specialization. We therefore adopt the Medical Abstracts dataset (Schopf et al., 2023), which consists of paper abstracts labeled by disease categories, simulating the task of narrow domain experts classifying diseases.

Both EDINET-Bench Extended and Medical Abstracts are imbalanced across classes. Nevertheless, our assumption in Eq. (1) is that complementary labels are sampled uniformly at random, independent of the original label distribution. We follow this procedure and empirically verify that class imbalance does not affect the validity of our estimator.

**Results** As shown in Table 2, both IVW and ML consistently achieve the lowest average per-run standard deviation, confirming the variance reduction effect observed in § 3. In contrast, the Comp-(K-1) estimator exhibits much larger variance on EDINET-Bench, which aligns with Eq. (4): with K=16 industries and relatively low accuracy, the variance is amplified.

Comparing EDINET-Bench and its extended version, we observe that increasing the number of samples substantially reduces the fluctuations of the estimators, supporting our earlier finding in § 3 that small-scale benchmarks such as GPQA can lead to unstable estimates. Finally, on the Medical Abstracts dataset, IVW and ML again achieve the most stable performance, demonstrating that the complementary-label estimators generalize across domains and languages.

**Agent Search with Estimator Feedback** Finally, we show that weak human supervision can be leveraged not only for evaluation but also as training signal for AI systems. We evaluate agentic flows under complementary-label signals on three benchmarks. GPQA serves as the most challenging setting in our study (Table 1); Math-MC focuses on reasoning ability; and Medical Abstracts best captures our practical scenario. Following Hu et al. (2025), we include CoT (Wei et al., 2022), CoT-SC (Wang et al., 2023), Self-Refine (Madaan et al., 2023), LLM Debate (Du et al., 2024), Step-back Abstraction (Zheng et al., 2024), Quality-Diversity (Lu et al., 2025), and Role Assignment (Xu et al., 2025) as manually designed baselines, and compare against ADAS (Hu et al., 2025) and AFlow (Zhang et al., 2025c). We define the accuracy computed by *q* as *complementary accuracy* and apply Eq. (2) to obtain the *transformed complementary accuracy*. For ADAS and AFlow, we follow the original protocol (see Appendix J), computing both complementary and transformed complementary accuracy on the validation set and reporting the better-performing variant.

Figure 3 compares ADAS and AFlow against strong baselines on GPQA, Math-MC, and Medical Abstracts. Both agentic systems consistently outperform manually designed workflows once weak complementary-label signals are incorporated, demonstrating that weak human feedback can provide learning signal. While the choice of estimator (raw vs. transformed complementary accuracy)

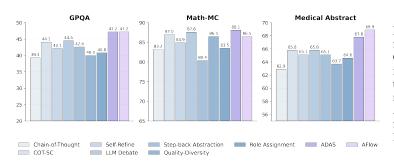


Figure 3: Accuracy across different benchmarks. Across GPQA, Math-MC, and Medical Abstract, agentic systems guided by weak complementary signals (ADAS and AFlow) consistently outperform manually designed baselines.

leads to different outcomes depending on task difficulty, the overall pattern remains robust: weak signals suffice to improve performance beyond established baselines, with additional variance analysis provided in Appendix K. These findings reinforce the central theme of our study: humans' weak labels can provide a reliable and scalable supervisory channel for steering agentic AI systems.

#### 4 RELATED WORK

Scalable oversight In superhuman regimes where AI systems surpass human abilities, humans can no longer provide reliable supervision to evaluate and train them. Recent scalable-oversight work explores strategies that let limited (human) supervisors steer stronger models. Among the many scalable oversight protocols (such as weak-to-strong generalization (Burns et al., 2024), easy-to-hard generalization (Sun et al., 2024), debate (Khan et al., 2024), and RLAIF (Bai et al., 2022)), our approach may seem similar to iterated amplification/IDA (Christiano et al., 2018) and recursive task decomposition (Wu et al., 2021) at a glance. The difference is that our approach can be used for tasks that are not easily decomposable, as we are not decomposing the task, but rather relying on expertise-wise partitioned human supervision to enable the usage of complementary labels.

**Weak supervision** Weakly-supervised learning (Sugiyama et al., 2022) studies learning with weaker form of supervision instead of ground-truth labels. Complementary labels (Ishida et al., 2017; 2019; Yu et al., 2018; Wang et al., 2024a; 2025), in particular, provide "not-this-class" information, with the traditional motivation of reducing annotation cost. Our perspective differs: we leverage weak supervision not merely for cheaper labeling, but as a principled route to scalable oversight in the superhuman regime, where no single human expert can provide verification end-to-end.

**Agentic systems** LLMs have inspired a wide range of system architectures (Wei et al., 2022; Wang et al., 2023; Madaan et al., 2023; Shinn et al., 2023b;a), in particular by incorporating multiple LLMs into collaborative frameworks (Du et al., 2024; Liu et al., 2024; Chen et al., 2024; Wu et al., 2024; Wang et al., 2024c). Early studies relied on manually designed agentic workflows, which demanded substantial manual design and domain expertise. More recent work has shifted toward *automated agent discovery*, in which LLMs themselves are used to generate, refine, and optimize workflow structures (Hu et al., 2025; Zhang et al., 2025a;b;c), but required ground-truth supervision to construct a fitness score. We investigate whether automated agentic system design can proceed even with weaker training signals, i.e., complementary labels.

#### 5 CONCLUSION

We introduced *scalable oversight via partitioned human supervision*, a protocol that utilizes the weak human supervision for tasks beyond any single expert's ability. We derived an unbiased estimator of performance metric from complementary labels, analyzed its variance, and provided finite-sample guarantees. We further proposed two practical mixture estimators that combine scarce ordinary labels with abundant complementary labels. Empirically, we show that these estimators allow us to estimate the performance of AI systems without requiring full ground truth. They also serve as effective training signals which we demonstrate with automatically designing agentic workflows.

### ETHICS STATEMENT

Our work is motivated by AI safety: we aim to improve oversight of advanced AI systems by enabling evaluation and training even on tasks beyond any single human expert's competence. This approach is intended to enhance the reliability and alignment of future AI models with human interests, which we believe is a positive contribution to society.

Currently, we do not foresee immediate direct negative social impacts from the research itself. However, we recognize that alignment techniques are a double-edged sword (Machine Learning Street Talk, 2025). While our framework is intended to improve safety, a malicious actor could in principle repurpose similar methods to align an AI system with harmful or unethical objectives. This risk is not unique to our approach, as any alignment method can be misused if directed toward bad ends. We explicitly caution against such misuse. Our work is presented in the spirit of enhancing beneficial oversight, and we emphasize that strong governance and ethical guidelines are needed to ensure alignment techniques are only applied to promote safety in AI systems (Bengio et al., 2024).

#### REPRODUCIBILITY STATEMENT

All code and scripts required to reproduce our experiments are included in the submission repository. For ADAS (Hu et al., 2025) and AFlow (Zhang et al., 2025c), we primarily follow the official public implementations; see Appendix J for additional details. Since our study relies on the OpenAI API, exact replication may be affected if certain models are deprecated or become unavailable in the future. Nevertheless, we provide complete implementation details, dataset references, and evaluation scripts to ensure reproducibility to the greatest extent possible.

#### Datasets.

- https://huggingface.co/datasets/TIGER-Lab/MMLU-Pro
- https://huggingface.co/datasets/GBaker/ MedQA-USMLE-4-options
- https://huggingface.co/datasets/Idavidrein/gpqa
- https://huggingface.co/datasets/SakanaAI/EDINET-Bench
- https://huggingface.co/datasets/TimSchopf/medical\_abstracts

#### Code.

- https://github.com/Geralt-Targaryen/MC-Evaluation
- https://github.com/SakanaAI/edinet2dataset
- https://github.com/ShengranHu/ADAS
- https://github.com/FoundationAgents/AFlow

#### REFERENCES

Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.

Anthropic. Introducing claude 4: Claude opus 4 and claude sonnet 4. Blog post on the Anthropic website, May 22 2025. URL https://www.anthropic.com/news/claude-4. Accessed 2025-09-24.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional AI: Harmlessness from AI feedback. *arXiv preprint arXiv:2212.08073*, 2022.

Yoshua Bengio, Geoffrey Hinton, Andrew Yao, Dawn Song, Pieter Abbeel, Trevor Darrell, Yuval Noah Harari, Ya-Qin Zhang, Lan Xue, Shai Shalev-Shwartz, et al. Managing extreme AI risks amid rapid progress. *Science*, 384(6698):842–845, 2024.

- Stéphane Boucheron, Gábor Lugosi, and Olivier Bousquet. Concentration inequalities. In *Summer school on machine learning*, pp. 208–240. Springer, 2003.
  - Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
  - Samuel R. Bowman, Jeeyoon Hyun, Ethan Perez, Edwin Chen, Craig Pettit, Scott Heiner, Kamilė 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.
  - Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeffrey Wu. Weak-to-Strong Generalization: Eliciting Strong Capabilities With Weak Supervision. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
  - Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *Proceedings of the Twelfth International Conference on Learning Representations*, 2024.
  - Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters, 2018.
  - Paul Christiano, Buck Shlegeris, and Dario Amodei. Supervising strong learners by amplifying weak experts. *arXiv preprint arXiv:1810.08575*, 2018.
  - Google DeepMind. Introducing gemini 2.5 pro and gemini 2.5 flash. Blog post on the Google DeepMind website, May 14 2025. URL https://deepmind.google/technologies/gemini/. Accessed 2025-09-24.
  - DeepSeek. Deepseek r1 updates: Build r1-0528. Model release notes on the DeepSeek platform, May 28 2025. URL https://platform.deepseek.com. Accessed 2025-09-24.
  - Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving Factuality and Reasoning in Language Models through Multiagent Debate. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
  - 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. In *Advances in Neural Information Processing Systems*, 2021.
  - Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
  - Takashi Ishida, Gang Niu, Weihua Hu, and Masashi Sugiyama. Learning from complementary labels. In *Advances in Neural Information Processing Systems*, 2017.
  - Takashi Ishida, Gang Niu, Aditya Menon, and Masashi Sugiyama. Complementary-label learning for arbitrary losses and models. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
  - Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14), 2021.

- Akbir Khan, John Hughes, Dan Valentine, Laura Ruis, Kshitij Sachan, Ansh Radhakrishnan, Edward Grefenstette, Samuel R. Bowman, Tim Rocktäschel, and Ethan Perez. Debating with More Persuasive LLMs Leads to More Truthful Answers. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
  - Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James Validad Miranda, Alisa Liu, Nouha Dziri, Xinxi Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Christopher Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing frontiers in open language model post-training. In *Proceedings of the Second Conference on Language Modeling*, 2025.
  - Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. A dynamic LLM-powered agent network for task-oriented agent collaboration. In *Proceedings of the First Conference on Language Modeling*, 2024.
  - Cong Lu, Shengran Hu, and Jeff Clune. Intelligent go-explore: Standing on the shoulders of giant foundation models. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
  - Machine Learning Street Talk. Don't invent faster horses Prof. Jeff Clune. YouTube video, January 2025. URL https://www.youtube.com/watch?v=mw5WIDGRLnA. Accessed: 2025-09-24.
  - Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*, 2023.
  - Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample-variance penalization. In *Proceedings of the 22nd Conference on Learning Theory*, 2009.
  - OpenAI. Introducing gpt-5. Blog post on the OpenAI website, August 7 2025a. URL https://openai.com/index/introducing-gpt-5/. Accessed 2025-09-24.
  - OpenAI. Introducing gpt-5 for developers: Gpt-5, gpt-5-mini, and gpt-5-nano. OpenAI Blog / API documentation, August 7 2025b. URL https://openai.com/index/introducing-gpt-5-for-developers. Includes the GPT-5 nano model; developer API pricing and performance trade-offs.
  - Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, 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*, 2022.
  - David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level google-proof q&a benchmark. In *Proceedings of the First Conference on Language Modeling*, 2024.
  - Tim Schopf, Daniel Braun, and Florian Matthes. Evaluating unsupervised text classification: Zeroshot and similarity-based approaches. In *Proceedings of the 6th International Conference on Natural Language Processing and Information Retrieval*, 2023.
  - Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
  - Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024.

- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, 2023a.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, 2023b.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. Learning to summarize with human feedback. In *Advances in Neural Information Processing Systems*, 2020.
- Issa Sugiura, Takashi Ishida, Taro Makino, Chieko Tazuke, Takanori Nakagawa, Kosuke Nakago, and David Ha. Edinet-bench: Evaluating llms on complex financial tasks using japanese financial statements, 2025.
- Masashi Sugiyama, Han Bao, Takashi Ishida, Nan Lu, and Tomoya Sakai. *Machine learning from weak supervision: An empirical risk minimization approach*. MIT Press, 2022.
- Zhiqing Sun, Longhui Yu, Yikang Shen, Weiyang Liu, Yiming Yang, Sean Welleck, and Chuang Gan. Easy-to-hard generalization: Scalable alignment beyond human supervision. *Advances in Neural Information Processing Systems*, 2024.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Hsiu-Hsuan Wang, Mai Tan Ha, Nai-Xuan Ye, Wei-I Lin, and Hsuan-Tien Lin. CLImage: Humanannotated datasets for complementary-label learning. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856.
- Wei Wang, Takashi Ishida, Yu-Jie Zhang, Gang Niu, and Masashi Sugiyama. Learning with complementary labels revisited: The selected-completely-at-random setting is more practical. In *Proceedings of the 41st International Conference on Machine Learning*, 2024a.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhu Chen. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, 2024b.
- Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei, and Heng Ji. Unleashing the emergent cognitive synergy in large language models: A task-solving agent through multi-persona self-collaboration. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2024c.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), Advances in Neural Information Processing Systems, 2022.
- Jeff Wu, Long Ouyang, Daniel M Ziegler, Nisan Stiennon, Ryan Lowe, Jan Leike, and Paul Christiano. Recursively summarizing books with human feedback. *arXiv preprint arXiv:2109.10862*, 2021.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. Autogen: Enabling next-gen LLM applications via multi-agent conversations. In *Proceedings of the First Conference on Language Modeling*, 2024.

- Stefan Wuchty, Benjamin F Jones, and Brian Uzzi. The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039, 2007.
- Benfeng Xu, An Yang, Junyang Lin, Quan Wang, Chang Zhou, Yongdong Zhang, and Zhendong Mao. Expertprompting: Instructing large language models to be distinguished experts, 2025.
- Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025.
- Xiyu Yu, Tongliang Liu, Mingming Gong, and Dacheng Tao. Learning with biased complementary labels. In *Proceedings of the European conference on computer vision*, 2018.
- Guibin Zhang, Luyang Niu, Junfeng Fang, Kun Wang, LEI BAI, and Xiang Wang. Multi-agent architecture search via agentic supernet. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025a.
- Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin godel machine: Openended evolution of self-improving agents. *arXiv* preprint arXiv:2505.22954, 2025b.
- Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xiong-Hui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, and Chenglin Wu. AFlow: Automating agentic workflow generation. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025c.
- Ziyin Zhang, Lizhen Xu, Zhaokun Jiang, Hongkun Hao, and Rui Wang. Multiple-choice questions are efficient and robust llm evaluators, 2024.
- Huaixiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed H. Chi, Quoc V Le, and Denny Zhou. Take a step back: Evoking reasoning via abstraction in large language models. In *Proceedings of the Twelfth International Conference on Learning Representations*, 2024.
- Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv* preprint arXiv:1909.08593, 2019.

### A PROOF OF COROLLARY 1

*Proof.* For a single item, under Eq. (1),

$$\begin{split} \mathbb{E}[W] &= \Pr(\hat{y} = y) \Pr(W = 1 \mid \hat{y} = y) + \Pr(\hat{y} \neq y) \Pr(W = 1 \mid \hat{y} \neq y) \\ &= \Pr(\hat{y} = y) \cdot 1 + \Pr(\hat{y} \neq y) \cdot \left(1 - \frac{1}{K - 1}\right) = \frac{A + K - 2}{K - 1}. \end{split}$$

If we consider  $\hat{q} = \frac{1}{n_c} \sum_{i=1}^{n_c} W_i$ , an empirical version of  $\mathbb{E}[W]$ , this  $\hat{q}$  is an unbiased estimator of  $\mathbb{E}[W]$ . Hence,  $\hat{A}_{\text{comp}} = (K-1)\hat{q} - (K-2)$  is an unbiased estimator of A.

# B CONNECTION BETWEEN THE COMPLEMENTARY-LABEL ESTIMATOR AND COMPLEMENTARY LOSS

**Notation.** Let  $(X,Y) \sim \mathcal{D}$  with  $Y \in [K] = \{1,\ldots,K\}$ . Given a model g, denote its prediction by  $\hat{y} = g(X)$ . For any class  $k \in [K]$ , the loss is  $\ell(k,\hat{y})$ . In complementary-label learning Ishida et al. (2019), one observes  $\bar{Y} \neq Y$  sampled uniformly from  $[K] \setminus \{Y\}$ . The complementary loss is defined (Eq. (10) in Ishida et al. (2019)) as

$$\bar{\ell}(k,\hat{y}) = -(K-1)\,\ell(k,\hat{y}) + \sum_{j=1}^{K}\ell(j,\hat{y}).$$

**Specialization to** 0-1 **loss.** For the 0–1 loss  $\ell(k, \hat{y}) = \mathbb{I}\{\hat{y} \neq k\}$ , we compute

$$\bar{\ell}(k,\hat{y}) = -(K-1)\,\mathbb{I}\{\hat{y} \neq k\} + \sum_{j=1}^{K} \mathbb{I}\{\hat{y} \neq j\}$$

$$= -(K-1)\,\mathbb{I}\{\hat{y} \neq k\} + (K-1)$$

$$= (K-1)\,\mathbb{I}\{\hat{y} = k\}.$$
(11)

Risk under complementary labels. The standard risk is

$$R = \mathbb{E}_{(X|Y)}[\ell(Y, \hat{y})] = \Pr(\hat{y} \neq Y) = 1 - A,$$

where  $A = \Pr(\hat{y} = Y)$  is the accuracy. By the risk-rewrite identity of Ishida et al. (2019) (Eq. (8)), we also have

$$R = \mathbb{E}_{(X,\bar{Y})}[\bar{\ell}(\bar{Y},\hat{y})] = (K-1) \ \mathrm{Pr}(\hat{y} = \bar{Y}).$$

**Equivalence to our estimator.** Let  $W = \mathbb{I}\{\hat{y} \neq \bar{Y}\}$ . Then  $\Pr(\hat{y} = \bar{Y}) = 1 - \mathbb{E}[W]$ , so

$$1 - A = (K - 1)(1 - \mathbb{E}[W]).$$

Rearranging yields

$$A = (K-1)\mathbb{E}[W] - (K-2).$$

This is exactly the linear correction in Eq. (2). Replacing  $\mathbb{E}[W]$  by its empirical mean  $\hat{q} = \frac{1}{n_c} \sum_{i=1}^{n_c} W_i$  gives the unbiased estimator  $\widehat{A}_{\text{comp}} = (K-1)\widehat{q} - (K-2)$  of Theorem 1.

#### C DATA LABELING AND ESTIMATION PROTOCOL

In multiple-choice evaluation, the effective "class" is the *index position* of an option rather than its textual content. To satisfy the uniformity condition in Eq. (1), the labeling process proceeds as follows. For each item, all options are first shuffled uniformly at random. After shuffling, the correct index y is recorded. A complementary label is then created by sampling a wrong index  $\bar{y}$  uniformly from the K-1 remaining positions. The final published datum is  $(x, \text{shuffled options}, \bar{y})$ . When the ground truth is known, this procedure can be automated: y is only used internally for generating  $\bar{y}$ , and the released dataset never exposes y. When labels are collected from human raters, the annotation interface can enforce the same randomization in the backend to guarantee uniformity.

811

812

813

814

815

816

817

818

819 820

821 822

823

824

826

827

828

829

830

831

832

835 836

837

838

839

840

841

842 843 844

845 846

847

848

849 850

851 852

853 854

855 856

857

858

859 860 861

862

863

Any deviation, such as asking raters to pick the "most plausible wrong option," would violate the assumption behind Eq 1.

The same protocol can be extended to cases where the ground truth is not available in advance. A complementary label is first generated in the same way, and then the annotator is asked to verify whether it happens to be correct. This verification yields a mixture of ordinary and complementary labels, which forms the mixed dataset studied in § 2.2.

The end-to-end evaluation pipeline, including both ordinary and complementary sets and the associated estimators, is summarized in Algorithm 1. Two estimation strategies are supported: a closed-form maximum-likelihood estimator and an inverse-variance weighted (IVW) estimator.

#### **Algorithm 1** Collecting complementary labels and estimating accuracy (fixed K)

```
1: Input: MCQ items; evaluation model f; sizes n_0, n_c.
            2: for each item do
            3:
                    Shuffle options uniformly; record correct index y.
            4:
                    if item \in ordinary set then
825
                         Evaluate f to obtain prediction \hat{y}; record Z = \mathbb{I}\{\hat{y} = y\}.
            5:
                                                                                                          6:
                         Sample \bar{y} uniformly from \{1, \dots, K\} \setminus \{y\}; publish (x, \text{options}, \bar{y}).
            7:
                        Evaluate f to obtain prediction \hat{y}; record W = \mathbb{I}\{\hat{y} \neq \bar{y}\}.
            8:
            9:
                    end if
           10: end for
           11: Compute S_0 = \sum Z, S_c = \sum W.
          12: Option A (ML): return \widehat{A}_{\rm ML} via Eq. (9) and a confidence interval from observed information.
833
           13: Option B (IVW): compute \widehat{A}_{\text{ord}}, \widehat{A}_{\text{comp}}, and \widehat{w}, then return \widehat{A}_{\text{IVW}} with plug-in confidence
834
               interval.
```

Two assumptions are critical for the validity of these estimators. First, the wrong index in each complementary label must be sampled uniformly; systematic biases, such as consistently preferring hard negatives, lead to biased estimates of  $\hat{A}_{\text{comp}}$ . Second, the ordinary and complementary sets must be drawn from the same distribution of items. If domain shift or subject imbalance exists, the combined estimator reflects accuracy on a mixture distribution. In such cases, stratified or subjectwise weighting may be preferable.

#### ONE-STEP NEWTON UPDATE AND IVW EQUIVALENCE D

**Setup and notation.** Let  $A \in (0,1)$  denote the (unknown) accuracy for a K-way multiple-choice task  $(K \geq 3)$ . We observe two independent samples: (i) ordinary labels  $S_o \sim \text{Bin}(n_o, A)$  with empirical accuracy  $A_{\rm ord} = S_{\rm o}/n_{\rm o}$ ; (ii) complementary labels obtained by uniformly sampling a wrong option per item, so that  $S_c \sim \text{Bin}(n_c, q)$  with

$$q(A) = \frac{A+K-2}{K-1}, \qquad \widehat{q} = S_{\rm c}/n_{\rm c}, \qquad \widehat{A}_{\rm comp} = (K-1)\widehat{q} - (K-2).$$

We consider a pilot  $A_0$  (e.g.  $\widehat{A}_{ord}$  or  $\widehat{A}_{comp}$ ) and write  $q_0 = q(A_0)$ .

Newton update and plug-in simplification. The joint log-likelihood is

$$\ell(A) = S_0 \log A + (n_0 - S_0) \log(1 - A) + S_c \log q + (n_c - S_c) \log(1 - q).$$

Differentiating gives the score functions

$$U_{\rm o}(A) = \frac{n_{\rm o}(\widehat{A}_{\rm ord} - A)}{A(1 - A)}, \qquad U_{\rm c}(A) = \frac{n_{\rm c}(\widehat{A}_{\rm comp} - A)}{(K - 1)^2 q(1 - q)}.$$

The observed information terms are

$$\mathcal{I}_{\rm o}^{\rm obs}(A) = \frac{S_{\rm o}}{A^2} + \frac{n_{\rm o} - S_{\rm o}}{(1 - A)^2}, \qquad \mathcal{I}_{\rm c}^{\rm obs}(A) = \frac{1}{(K - 1)^2} \left( \frac{S_{\rm c}}{q^2} + \frac{n_{\rm c} - S_{\rm c}}{(1 - q)^2} \right).$$

A one-step Newton update around  $A_0$  is

$$A_1 = A_0 - \frac{U(A_0)}{U'(A_0)} = A_0 + \frac{U_0(A_0) + U_c(A_0)}{\mathcal{I}_o^{\text{obs}}(A_0) + \mathcal{I}_c^{\text{obs}}(A_0)}.$$

When we set  $A_0=\widehat{A}_{\mathrm{ord}}$ , the ordinary part simplifies exactly: since  $S_{\mathrm{o}}=n_{\mathrm{o}}A_0$ , we have  $U_{\mathrm{o}}(A_0)=0$  and  $\mathcal{I}_{\mathrm{o}}^{\mathrm{obs}}(A_0)=n_{\mathrm{o}}/[A_0(1-A_0)]=I_{\mathrm{o}}(A_0)$ , the Fisher information. For the complementary part, we plug in  $q=\widehat{q}$  so that  $S_{\mathrm{c}}=n_{\mathrm{c}}\widehat{q}$  holds identically. This yields

$$\left. \mathcal{I}_{\mathrm{c}}^{\mathrm{obs}}(A_0) \right|_{q=\widehat{q}} = \frac{n_{\mathrm{c}}}{(K-1)^2 \, \widehat{q}(1-\widehat{q})} = I_{\mathrm{c}}(\widehat{q}), \qquad \left. U_{\mathrm{c}}(A_0) \right|_{q=\widehat{q}} = I_{\mathrm{c}}(\widehat{q}) \, (\widehat{A}_{\mathrm{comp}} - A_0).$$

Substituting back into the Newton update gives

$$A_1 \approx A_0 + \frac{I_c(\widehat{q})(\widehat{A}_{comp} - A_0)}{I_o(A_0) + I_c(\widehat{q})}.$$

**Closed-form IVW rule.** Rearranging shows that the practical one-step update is a convex combination of the two empirical estimators:

$$\widehat{A}_{\rm IVW} \; \approx \; \frac{I_{\rm o}(\widehat{A}_{\rm ord}) \, \widehat{A}_{\rm ord} + I_{\rm c}(\widehat{q}) \, \widehat{A}_{\rm comp}}{I_{\rm o}(\widehat{A}_{\rm ord}) + I_{\rm c}(\widehat{q})}. \label{eq:alpha}$$

Variance-information identities and final form. For the two constituent estimators, the variance-information identities are

$$\operatorname{Var}(\widehat{A}_{\operatorname{ord}}) = \frac{1}{I_o}, \quad \operatorname{Var}(\widehat{A}_{\operatorname{comp}}) = \frac{1}{I_c},$$

with the plug-in Fisher information 
$$I_{\rm o}(\widehat{A}_{\rm ord}) = \frac{n_{\rm o}}{\widehat{A}_{\rm ord}(1-\widehat{A}_{\rm ord})}, \qquad I_{\rm c}(\widehat{q}) = \frac{n_{\rm c}}{(K-1)^2\,\widehat{q}(1-\widehat{q})}.$$
 Therefore the IVW weight can be written either in information form,  $w^\star = \frac{I_{\rm o}}{I_{\rm ord}}$ .

Therefore the IVW weight can be written either in information form,  $w^* = \frac{I_0}{I_0 + I_c}$ , or, equivalently, in variance form (using  $I_{\bullet} = 1/\mathrm{Var}(A_{\bullet})$ ),

$$w^{\star} = \frac{I_{\rm o}}{I_{\rm o} + I_{\rm c}} = \frac{{\rm Var}(\widehat{A}_{\rm comp})}{{\rm Var}(\widehat{A}_{\rm ord}) + {\rm Var}(\widehat{A}_{\rm comp})}.$$

Plugging these into the convex combination yields exactly results in §2.2:

$$\widehat{A}_{\text{IVW}} = w^{\star} \widehat{A}_{\text{ord}} + (1 - w^{\star}) \widehat{A}_{\text{comp}}.$$

In summary, the IVW estimator is a one-step Newton approximation to the joint ML: it is first-order equivalent and asymptotically efficient, but may differ in finite samples because the weights are fixed by plug-in information rather than optimized jointly. While the ML is theoretically optimal, the IVW rule is a simple linear combination that often delivers practically comparable performance and is easier to apply in practice.

#### Ε PROOF OF THEOREM 2

*Proof.* Hoeffding part. For i.i.d.  $W_i \in [0, 1]$ , Hoeffding gives

$$\Pr(|\hat{q} - \mathbb{E}\hat{q}| \ge \epsilon) \le 2\exp(-2n_c\epsilon^2).$$

Since 
$$\widehat{A}_{\text{comp}} = (K-1)\widehat{q} - (K-2)$$
 and  $A = (K-1)\mathbb{E}\widehat{q} - (K-2)$ , we have

$$|\widehat{A}_{\text{comp}} - A| = (K - 1) |\widehat{q} - \mathbb{E}\widehat{q}|.$$

Empirical Bernstein part. Apply the two-sided empirical Bernstein inequality (Maurer & Pontil, 2009) to  $\hat{q} = \frac{1}{n_o} \sum_{i=1}^{n_c} W_i$  with  $W_i \in [0, 1]$ :

$$|\hat{q} - \mathbb{E}\hat{q}| \le \sqrt{\frac{2\,\hat{q}(1-\hat{q})\,\log(4/\delta)}{n_{\mathrm{c}}}} + \frac{7\,\log(4/\delta)}{3\,(n_{\mathrm{c}}-1)}.$$

Multiplying both sides by (K-1) and using  $\widehat{A}_{\text{comp}} - (A) = (K-1)(\widehat{q} - \mathbb{E}\widehat{q})$  yields the second branch in Theorem 2

#### F PROOF OF THEOREM 4

*Proof.* Recall  $Z_j = \mathbb{I}\{\hat{y}_j = y_j\}$ ,  $j = 1, \ldots, n_o$ , are i.i.d. Bernoulli(A) with  $\mathbb{E}[Z_j] = A$ ; and  $W_i = \mathbb{I}\{\hat{y}_i \neq \bar{y}_i\}$ ,  $i = 1, \ldots, n_c$ , are i.i.d. Bernoulli(q) with q = (A + K - 2)/(K - 1) under Eq. (1). The two groups are independent.

Define the rescaled variables

$$X_j^{(o)} := \frac{w}{n_o} Z_j, \qquad X_i^{(c)} := \frac{(1-w)(K-1)}{n_c} W_i,$$

and collect them as  $X_1, \ldots, X_n$  with  $n = n_0 + n_c$ . Then the centered sum

$$S := \sum_{j=1}^{n_o} (X_j^{(o)} - \mathbb{E}X_j^{(o)}) + \sum_{i=1}^{n_c} (X_i^{(c)} - \mathbb{E}X_i^{(c)}) = \widehat{A}_{\text{mix}} - A.$$
 (12)

We verify the conditions of Theorem 2.10 in Boucheron et al. (2003):

- (i) Independence) holds by construction.
- (ii) Boundedness). For all j, i, j

$$|X_j^{(o)} - \mathbb{E}X_j^{(o)}| \le \frac{w}{n_o}, \qquad |X_i^{(c)} - \mathbb{E}X_i^{(c)}| \le \frac{(1-w)(K-1)}{n_c}.$$

Set

$$c := \max \left\{ \frac{w}{n_o}, \frac{(1-w)(K-1)}{n_c} \right\}. \tag{13}$$

(iii) Second moment). Summing variances,

$$v := \sum_{j=1}^{n_{\rm o}} \text{Var}(X_j^{(o)}) + \sum_{i=1}^{n_{\rm c}} \text{Var}(X_i^{(c)}) = w^2 \frac{A(1-A)}{n_{\rm o}} + (1-w)^2 \frac{(K-1)^2 q(1-q)}{n_{\rm c}}.$$
 (14)

(iv) Higher moments). Since  $|X_i - \mathbb{E}X_i| \le c$ , for all integers  $q \ge 3$ ,  $(X_i - \mathbb{E}X_i)_+^q \le c^{q-2}(X_i - \mathbb{E}X_i)^2$ ; summing and taking expectations yields  $\sum_i \mathbb{E}[(X_i)_+^q] \le v \, c^{q-2}$ , which is dominated by the requirement  $\frac{q!}{2} \, v \, c^{q-2}$  in (Boucheron et al., 2003, Thm. 2.10). Therefore, Theorem 2.10 applies to S in Eq. (12) and (using the symmetric lower tail plus a union bound) gives  $\Pr\{|S| \ge \sqrt{2vt} + ct\} \le 2e^{-t}$ . Setting  $t = \log(2/\delta)$  yields Eq. (10).

#### G PRACTICAL PLUG-IN FORMS OF THE BERNSTEIN BOUND

The quantities A and q are unknown; replace them by their empirical counterparts. Define  $\hat{v}_{\rm o} = \frac{\hat{A}_{\rm ord}(1-\hat{A}_{\rm ord})}{n_{\rm o}}, \qquad \hat{v}_{\rm c} = \frac{(K-1)^2\,\hat{q}(1-\hat{q})}{n_{\rm c}}, \qquad \hat{v}_{\rm mix} = w^2\hat{v}_{\rm o} + (1-w)^2\hat{v}_{\rm c}, \text{ with } \hat{q} = S_{\rm c}/n_{\rm c}.$  Substituting into Eq. (10) yields the computable bound  $\left|\hat{A}_{\rm mix}-A\right| \leq \sqrt{2\,\hat{v}_{\rm mix}\,\log\frac{2}{\delta}} + c\,\log\frac{2}{\delta},$  where c is as in Eq. (13). For clarity, an equivalent compact statement is

$$|\hat{A}_{\text{mix}} - A| \le \sqrt{2\log\frac{2}{\delta}\left[w^2\hat{v}_0 + (1-w)^2\hat{v}_c\right]} + \log\frac{2}{\delta}\max\left(\frac{w}{n_o}, \frac{(1-w)(K-1)}{n_c}\right).$$
 (15)

The variance component is minimized at  $w^\star = \frac{v_{\rm c}}{v_{\rm o} + v_{\rm c}}$ , a practical w is defined as  $\hat{w}_{\rm IVW} = \frac{\hat{v}_{\rm c}}{\hat{v}_{\rm o} + \hat{v}_{\rm c}}$ , coinciding with the IVW weight in Eq. (7). At this choice, the dominant (variance) term equals  $\sqrt{2\log\frac{2}{\delta}\cdot\frac{v_{\rm o}v_{\rm c}}{v_{\rm o}+v_{\rm c}}}$ .

Defining  $\widehat{h}_{\min}:=\frac{\widehat{v}_o\,\widehat{v}_c}{\widehat{v}_o+\widehat{v}_c}$ , using the fully plug-in version gives

$$|\hat{A}_{\text{mix}} - A| \leq \sqrt{2\log\frac{2}{\delta}\,\hat{h}_{\text{min}}} + \log\frac{2}{\delta}\,\max\left(\frac{\hat{w}_{\text{IVW}}}{n_{\text{o}}},\,\frac{(1-\hat{w}_{\text{IVW}})(K-1)}{n_{\text{c}}}\right),\tag{16}$$

Conceptually, the mgf-based Bernstein bound (Eq. (10)) tightens Hoeffding's inequality by adding both a variance term v and a range-driven linear term c. As  $n_o$  or  $n_c$  grows, the right-hand

side vanishes: the variance part decays as  $O(1/n_{\rm o}) + O(1/n_{\rm c})$ , while the linear term scales as  $O(\max\{w/n_{\rm o},\,(1-w)(K-1)/n_{\rm c}\})$ . The plug-in IVW weight  $\widehat{w}_{\rm IVW}=\frac{\widehat{v}_{\rm c}}{\widehat{v}_{\rm o}+\widehat{v}_{\rm c}}$  converges to  $w^\star=\frac{v_{\rm c}}{v_{\rm o}+v_{\rm c}}$ ; in particular,  $w^\star\to 0$  when  $n_{\rm c}\to\infty$  with  $n_{\rm o}$  fixed (the complementary arm dominates), and  $w^\star\to 1$  when  $n_{\rm o}\to\infty$  with  $n_{\rm c}$  fixed (the ordinary arm dominates).

Formally, Eq. (10) is valid for any fixed choice of the weight  $w \in [0,1]$  (e.g., a constant choice or an estimate obtained from an independent pilot split). If instead w is chosen adaptively from the same evaluation data (e.g., the plug-in  $\widehat{w}_{\text{IVW}}$ ), the fixed-weight assumption is violated and the bound is no longer guaranteed. In such cases one can (i) revert to the union-bound inequalities of the main text, which hold for arbitrary (possibly data-dependent) w via a triangle inequality; (ii) restore the fixed-w condition through sample splitting or cross-fitting (Chernozhukov et al., 2018; Wager & Athey, 2018); or (iii) restrict w to a finite grid  $G \subset [0,1]$  and apply a union bound with  $\delta$  adjusted by |G| (Boucheron et al., 2013; Shalev-Shwartz & Ben-David, 2014).

#### H BENCHMARK DETAILS

From MMLU-Pro (Wang et al., 2024b), we select only the 10-option questions, yielding 9,795 examples. MedQA-USMLE (Jin et al., 2021) and GPQA (Rein et al., 2024) are both four-option multiple-choice benchmarks; for GPQA we adopt the extended version for broader coverage. For MATH-MC, following the candidate-generation protocol of Zhang et al. (2024), we construct five-option questions from the original MATH dataset (Hendrycks et al., 2021), discarding problems with fewer than four distractors, which yields 11,751 examples. All duplicates are removed to ensure consistency.

#### I ESTIMATOR CONFIGURATIONS

For completeness, we summarize the exact configurations of all estimators used in Table 1.

- Ord Ordinary-label estimator based on  $n_{\rm o}=300$  samples per dataset ( $n_{\rm o}=120$  for GPQA due to dataset size). Accuracy is computed directly from the proportion of correct ordinary labels.
- Comp- $n_{\rm o}$  Complementary-label estimator with the same number of samples as the ordinary case, i.e.  $n_{\rm c}=n_{\rm o}$  ( $n_{\rm c}=120$  for GPQA). This setting highlights the loss of efficiency when only complementary labels are available.
- Comp-(K-1) Complementary-label estimator with sample size fixed at (K-1) times the ordinary case, i.e.  $n_{\rm c}=(K-1)n_{\rm o}$ . Following the protocol in Appendix C, each instance with oracle label y is converted into (K-1) complementary "No" labels  $\{\bar{y}\in\mathcal{Y}\setminus\{y\}\}$ , and the ordinary label is discarded; thus no ordinary labels are present in this setting.
  - Comp-Var Complementary-label estimator with sample size  $n_{\rm c}$  chosen such that its variance matches that of  $n_{\rm o}$  ordinary labels, following Eq. (4). This represents the idealized regime where complementary labels are sufficiently abundant to offset their lower per-sample information.
    - IVW-0.5 A baseline inverse-variance weighted estimator with a fixed equal weight w=0.5 between ordinary and complementary accuracies. This serves as a naïve combination strategy without using variance information.
      - IVW Inverse-variance weighted estimator with closed-form plug-in weights (Eq. (6)), combining ordinary and complementary estimates optimally under the variance–information identities.
      - ML Maximum-likelihood estimator obtained by solving the joint likelihood equation using both ordinary and complementary samples. Asymptotically equivalent to IVW but involving a nonlinear update (Appendix D).
    - Ord-Eval Oracle reference: accuracy evaluated directly on the entire dataset using ordinary labels.

      This is shown only for comparison and is excluded from macro-averages.
    - **Notes on datasets.** † Math denotes Math-MC evaluated with GPT-4.1-nano instead of GPT-5-nano, due to reasoning-output constraints. ‡ Math (CoT) denotes the

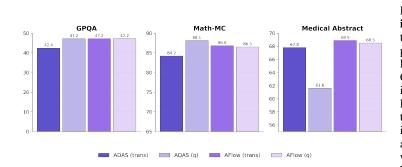


Figure 4: Estimator sensitivity under weak signals. Detailed comparison of raw complementary accuracy (q) and its linear transform (trans) across GPQA, Math-MC, and Medical Abstract. The figure illustrates how the two estimators behave differently depending on task difficulty, with variance amplification explaining when the transform helps or burts

same **Math-MC** dataset evaluated under chain-of-thought prompting (applied only to Math; model=GPT-4.1-nano).

## J IMPLEMENTATION DETAILS ON ADAS AND AFLOW

To ensure a fair comparison, we constructed validation and test splits for each benchmark, following the numbers and ratios used in Hu et al. (2025). Specifically, we sampled 120 validation and 800 test examples for Math-MC and Medical Abstract, and 88 validation and 458 test examples for GPQA. For AFlow, we set the number of validation runs to 3 and used the same validation set as ADAS, rather than the separate validation subset reported in Zhang et al. (2025c).

For code and prompts, we largely follow the official implementations of Hu et al. (2025) and Zhang et al. (2025c), with two minor modifications for reproducibility in our setting.

**ADAS.** The set of experts was adapted to better align with our benchmarks (e.g., replacing domain-specific experts in LLM Debate with ones relevant to GPQA and Medical Abstract). All other prompt components remain unchanged. We uploaded the full text of the adapted prompts and scripts in the supplementary material.

**AFlow.** Since the original AFlow does not directly correspond to the multiple-choice format, we modified the answer extraction and normalization functions, and used the following prompt:

```
Answer the following multiple choice question.
      {problem['Question']}
1061
      (A) {problem['A']}
       (B) {problem['B']}
1062
      (C) {problem['C']}
1063
      (D) {problem['D']}
1064
      At the very end, output the final answer in the format:
      ### X
1067
1068
      Where X is exactly one of A, B, C, or D.
1069
1070
      Do not add anything after this line. The last line of your output must
1071
      only contain '### X'. If you write anything after the line '### X', your
      answer will be marked as incorrect.
1072
```

#### K ESTIMATOR VARIANCE ANALYSIS

Figure 4 reports ADAS and AFlow under the two estimators—raw complementary accuracy (q) and its linear transform (trans)—each selected at the iteration that maximizes validation accuracy. The results show clear regime dependence: on **GPQA** (the hardest benchmark), q is more reliable (e.g., ADAS drops from 47.2 with q to 42.4 with trans, while AFlow remains unchanged at 47.2).

On **Math-MC**, ADAS still prefers q (88.1 vs. 84.2), whereas AFlow shows only a slight gain with trans (86.8 vs. 86.5). On **Medical Abstract**, where complementary accuracy is relatively high, the transform becomes beneficial: ADAS improves from 61.6 to 67.8, and AFlow from 68.5 to 68.9.

This behavior is explained by variance amplification: according to Eq. (3), the variance of the transformed estimator scales by  $(K-1)^2$  compared to q. When signals are weak (e.g., GPQA), variance amplification exacerbates noise and harms performance. In settings with stronger signals (Math-MC, Medical Abstract), the effect is less detrimental, and the transformed estimator occasionally yields modest gains. A systematic study of this trade-off is left for future work.

#### L LARGE LANGUAGE MODELS USAGE STATEMENT

LLMs assisted us in formulating our ideas, working on extensions and analysis, and polishing the writing. All of the LLM outputs were subsequently verified, revised and modified by the authors. Any errors or inappropriate statements remain the responsibility of the authors.