

Selective Annotation via Data Allocation: These Data Should Be Triageed to Experts for Annotation Rather Than the Model

Chen Huang^{♠♥}, Yang Deng[♣], Wenqiang Lei^{♠♥*}, Jiancheng Lv^{♠♥}, Ido Dagan[◇]

♠ Sichuan University ♣ Singapore Management University ◇ Bar-Ilan University

♥ Engineering Research Center of Machine Learning and Industry Intelligence,

Ministry of Education, China

huangc.scu@gmail.com ydeng@smu.edu.sg

{wenqianglei, lvjiancheng}@scu.edu.cn dagan@cs.biu.ac.il

Abstract

To obtain high-quality annotations under limited budget, semi-automatic annotation methods are commonly used, where a portion of the data is annotated by experts and a model is then trained to complete the annotations for the remaining data. However, these methods mainly focus on selecting informative data for expert annotations to improve the model predictive ability (i.e., triage-to-human data), while the rest of the data is indiscriminately assigned to model annotation (i.e., triage-to-model data). This may lead to inefficiencies in budget allocation for annotations, as easy data that the model could accurately annotate may be unnecessarily assigned to the expert, and hard data may be misclassified by the model. As a result, the overall annotation quality may be compromised. To address this issue, we propose a selective annotation framework called SANT. It effectively takes advantage of both the triage-to-human and triage-to-model data through the proposed error-aware triage and bi-weighting mechanisms. As such, informative or hard data is assigned to the expert for annotation, while easy data is handled by the model. Experimental results show that SANT consistently outperforms other baselines, leading to higher-quality annotation through its proper allocation of data to both expert and model workers. We provide pioneering work on data annotation within budget constraints, establishing a landmark for future triage-based annotation studies.

1 Introduction

Creating a high-quality and fully annotated dataset by human experts is criticized as being expensive. It is impossible to complete the annotation if we are operating on a limited budget (Chen et al., 2021), especially when dealing with large-scale data corpora on the web (Hedderich et al., 2021; Feng and Lapata, 2008). In this paper, we study the data

annotation problem within the constraints of limited budgets, where the goal is to achieve high-quality annotations despite having insufficient budgets to hire experts for annotating the entire dataset. To tackle this challenge, a typical solution is to conduct semi-automatic annotation (Chen et al., 2020b; Yao et al., 2023) (cf. Fig. 1 (middle)). It typically allocates the budget to hire experts for annotating a portion of the data, and then utilizes the annotated data to train a prediction model for completing the annotations for the remaining data. The choice of the model can vary, ranging from sophisticated models with large parameters (Bryant et al., 2017; Schulz et al., 2019; Hedderich et al., 2021) to lightweight models (Chen et al., 2020a; Desmond et al., 2021a) depending on the available hardware budgets. However, such methods often overlook the data allocation problem, i.e., the process of deciding which data should be annotated by human experts and which should be left for the prediction models to annotate. This may lead to inefficiencies in budget allocation for annotations, as certain data that the model could accurately annotate is unnecessarily assigned to experts. Consequently, the overall annotation quality within the same budget may be compromised.

In the light of this challenge, we take the first step to re-formulate the semi-automatic annotation under limited budgets as a data triage problem, which aims to optimize the annotation quality under limited budgets by determining the allocation of data to human annotator (namely triage-to-human data) and data to be assigned for model prediction (namely triage-to-model data). Existing semi-automatic annotation methods (Desmond et al., 2021a; Su et al., 2022; Li et al., 2023) mainly focus on selecting unlabeled data for expert annotations to improve the model predictive ability (i.e., triage-to-human data) through random strategies or active learning (AL), while the rest of data is indiscriminately assigned to model annotation (i.e.,

*Corresponding author.

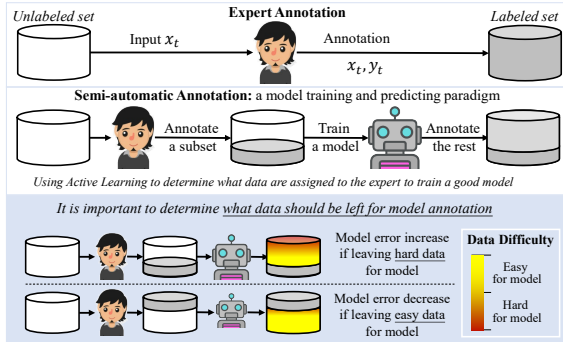


Figure 1: Semi-automatic annotation focuses on triage-to-human data, overlooking triage-to-model data.

trriage-to-model data). However, as shown in Fig. 1 (bottom), even if we constantly improve the predictive ability via triage-to-human data, the model still inevitably makes errors in its annotations. Not to mention that the data allocation strategies of existing methods may not effectively improve the model predictive ability (Tang and Huang, 2019; Mindermann et al., 2022). To address this issue, determining what data should be assigned for model annotation (trriage-to-model data) is crucial. A model could achieve high-quality annotations if assigned with the data that is easy for model to predict, while experts annotate the hard ones. Therefore, when solving the data triage challenge, it’s important to take both triage-to-human and triage-to-model data into account. In this way, the budget utilization could be improved through efficiently allocating both human and model annotations.

To this end, we propose a novel Selective ANnoTation framework (SANT), which optimizes the utilization of limited annotation budgets by balancing the importance of triage-to-human and triage-to-model data. Firstly, an AL-based mechanism is deployed to assign higher weights to data that help improve the model predictive ability. Secondly, an Error-Aware Triage (EAT) mechanism is proposed to assess the "hardness" weight of data, with the aim of redirecting the focus toward triage-to-model data. In specific, EAT dynamically estimates the probability of model error using triage-to-human data from previous rounds. Finally, a bi-weighting mechanism is adopted to aggregate different advantages of prioritizing triage-to-human and triage-to-model data. Such mechanism adaptively adjusts the importance of data so that informative or hard-to-predict data is assigned to humans for annotation. As such, the annotation budget are efficiently utilized to secure the model predictive ability and to leverage the strengths of the expert on hard data annotation simultaneously.

Our experiment results show that SANT consistently improves annotation quality across various annotation budgets, surpassing the best semi-automatic annotation method by an average of +0.50%, +4.86%, and +4.54% on the three annotation tasks, respectively. Recently, as large language models (LLMs) possess strong zero-shot and few-shot capabilities for data annotation (He et al., 2023; Gilardi et al., 2023; Ding et al., 2023), we further adopt ChatGPT as a strong automatic annotation baseline. Notably, SANT exhibits a better annotation quality over ChatGPT, even with the aid of few-shot in-context learning and Chain-of-Thought (CoT) prompting (Brown et al., 2020; Wei et al., 2022). These results highlight that proper allocation of data to both human and model workers results in higher-quality annotation. Further analysis highlights the importance of the triage-to-model data, which are more beneficial to obtain higher-quality annotations compared to the triage-to-human ones. In conclusion, we pioneer work on data annotation with limited budgets by formulating a data triage problem. We set a landmark for future triage-based annotation methods. Our contributions are as follows:

- We present the first work to study the roles of both triage-to-human and triage-to-model data. We verify the importance of triage-to-model data, which is largely overlooked by existing studies.
- We re-formulate semi-automatic annotation as data triage and propose a selective annotation framework, achieved by optimizing two novel mechanisms (i.e., EAT and bi-weighting mechanisms) for utilizing limited budgets.
- We show that proper allocation of data to both the expert and model workers results in higher-quality annotation. SANT effectively balances the advantages of two factors: improving model predictive ability (i.e., *trriage-to-human data*) and triaging data based on the model’s predictive ability (*trriage-to-model data*).

2 Related Work

Data Annotation. Human annotations are known to be expensive (Chen et al., 2021; Hedderich et al., 2021; Feng and Lapata, 2008). To this end, many studies have explored human-machine cooperative data annotation (Lu et al., 2024; Li et al., 2023; Ding et al., 2024; Huang et al., 2024a,b). In this case, a portion of the data is designated for human annotation (trriage-to-human data), while

the remaining data is allocated for model annotation (triage-to-model data). To further enhance the model’s predictive capabilities, existing methods are often formulated as a "model training and prediction" paradigm, where the model iteratively learns from the annotated data and then makes predictions. (Dalvi et al., 2016; Li et al., 2015; Bryant et al., 2017; Chen et al., 2020b; Wang et al., 2023; Tan et al., 2023). Considering the hardware budgets and time efficiency requirements, light-weighted models (Desmond et al., 2021a; Chen et al., 2020a), such as label propagation (Zhou et al., 2003) and MLP are preferred. Recently, uncertainty-based methods have been used to select a data subset for human annotations (Zhang et al., 2021; Desmond et al., 2021a; Su et al., 2022; Ein-Dor et al., 2020), improving the overall annotation quality by refining the model using more informative data. However, current methods concentrate solely on selecting triage-to-human data, while neglecting the importance of triage-to-model data. This fails to optimize data allocation effectively. In contrast, SANT considers both types of data.

Active learning. Current data allocation practices in the field of data annotation draw inspiration from the active learning (Ren et al., 2020), where data with high uncertainty are prioritized for human annotation. This aligns with active learning’s preference for informative or representative data (Desmond et al., 2021b; Loquercio et al., 2020). To our knowledge, no existing research has explored the application of other active learning techniques for data allocation in annotation. Our work, SANT, extends this approach by introducing a novel perspective. It prioritizes triage-to-model data, essentially introducing the anti-current learning (Braun et al., 2017) into the task of data allocation (i.e., we assign hard-to-predict data for human annotation). SANT seamlessly integrates active learning algorithms with the anti-current learning approach through our proposed bi-weighting mechanism. Our experiments demonstrate the effectiveness of this novel framework in enhancing the quality of data annotation.

3 SANT: Selective Annotation

We introduce the proposed Selective Annotation (SANT) framework. In Section 3.1, we provide the problem definition of the data annotation process, while Section 3.2 introduces the proposed framework SANT and two novel mechanisms. In

Appendix F, we introduce the learning details of our framework. Following previous works on semi-automatic annotation (Chen et al., 2020a; Desmond et al., 2021a; Hedderich et al., 2021) and human-model interaction (Hwa, 2000; Kristjansson et al., 2004; Huang et al., 2023), we assume the cost required for human annotation is constant for each data point, and estimate the annotation budgets using the number of data points.

3.1 Problem Definition

We reformulate the data annotation under limited budgets as a data triage problem, where a data triage module d assigns data to either the expert or the model f on-the-fly. The triage module outputs a triage signal of either 0 or 1 to delegate the annotation duty. If the signal is 0, the model is responsible for data annotation. Conversely, if the signal is 1, the human is assigned, and the model is updated using previous human-annotated data.

The annotation process terminates when one of two conditions is met: either the limited budgets are exhausted or all the data is annotated. In the first case, if the budgets are depleted before completing the entire dataset, the model annotator must continue with the remaining annotations. In the second case, it is possible for all data to be annotated before the budgets are exhausted if the amount of data assigned to the human during process is less than the budgets. To maximize the utilization of the human budget, a portion of the model-annotated data, selected by the customized data triage strategy, is re-allocated to the human annotator until their budget is depleted.

3.2 Proposed Framework

Semi-automatic annotation methods use active learning (AL) in the data triage module, which primarily prioritize triage-to-human data but overlook the triage-to-model data. As a result, they fail to efficiently and directly optimize budget utilization. To this end, SANT introduces two innovative mechanisms, namely EAT and bi-weighting mechanisms, which consider both human and model-based triage data for budget optimization.

3.2.1 AL-based mechanism

Semi-automatic annotation and SANT are flexible to any off-the-shelf AL methods¹, such as,

¹The general preference is for stream-based AL methods over pool-based ones, as they tend to be more time-efficient.

leveraging the uncertainty of the model prediction (Desmond et al., 2021b) or using a differentiable neural network (Wang et al., 2022). Formally speaking, given data x_t and the model f_t at round t , x_t is assigned to the expert if the AL-based mechanism d_t^{AL} gives x_t a large weight. If the AL-based mechanism is implemented using a learnable neural network, both AL and f_t are jointly optimized/updated, otherwise only f_t is updated. Note that the AL-based mechanism solely focuses on triage-to-human data, which limits its ability to optimize budget utilization directly and efficiently.

3.2.2 Error-aware triage mechanism (EAT)

Unlike existing AL-based semi-automatic annotation, we propose EAT to prioritize triage-to-model data by assigning hard-to-predict data to the expert and reserving easy-to-predict data to the model. Specifically, EAT assigns data x_t to the expert if the EAT-based mechanism d_t^{EAT} gives x_t a high weight, which is determined by the probability of f_t misclassifying x_t . The model f_t and EAT are then jointly optimized/updated. If d_t^{EAT} does not assign a high weight to x_t , then the model f_t acts as the annotator without requiring human involvement. Notably, EAT can be parameterized using any neural network, depending on the available hardware budgets. Here, EAT’s input and learning objectives are presented below.

Input of EAT. Instead of learning from training data independently (Mozannar and Sontag, 2020; Wilder et al., 2021; Raghu et al., 2019; Okati et al., 2021), to enhance EAT’s input features, we propose capturing the relationship between input data and its neighboring data from the same batch. This includes the use of weighted neighborhood entropy, which measures the model’s predictive uncertainty on the local area of input data x_t . To calculate this entropy, we retrieve the top- k neighbors $N_t(x_t)$ and calculate the cosine similarity between x_t and each neighbor $x_i \in N_t$. We store these similarity values in vector $C_t(x_t) \in R^k$. Then, we calculate the entropy of the model prediction $f_t(x_i)$ for each $x_i \in N_t(x_t)$ and store them in vector $E_t(x_t) \in R^k$. Finally, we obtain the weighted neighborhood entropy of x_t by using the element-wise multiplication operator \odot on vectors $E_t(x_t)$ and $C_t(x_t)$. The input of EAT is then the concatenation of the embedding of unlabeled data x_t , the model output $f_t(x_t)$, and the weighted neighborhood entropy.

Learning objective of EAT. EAT is designed to allocate hard-to-predict data, which are beyond the

model’s predictive ability, to the expert. To achieve this, EAT estimates the likelihood of the model making errors on unlabeled data x_t . Specifically, EAT is required to minimize the following loss.

$$L_d = \sum_{\{x_t, y_t\}} \ell_d(\mathbb{1}[y_t \neq \max(f_t(x_t))], d_t^{EAT}(x_t)). \quad (1)$$

Here, $\mathbb{1}[\cdot]$ is an indicator function. $y_t \in R$ and $f_t(x_t) \in R^C$ are the ground truth from the human and predicted outputs of the model, where C is the number of classes. Also, the operator $\max(\cdot)$ outputs the index of the maximal value of the input vector. In our experiments, ℓ_d comes in the form of the negative log-likelihood or binary cross-entropy loss, depending on the annotation tasks. Moreover, We further follow recent works on few-shot learning (Wang et al., 2018; Li et al., 2020) and utilize the max-margin loss to improve SANT’s ability to learn and generalize from limited labeled data on the fly. Apart from the aforementioned loss L_d , we add the following max-margin loss L_m (See Appendix C for ablation experiments on L_m).

$$\begin{aligned} L_m &= \max(0, \alpha + \overline{L}_f^m - \overline{L}_f^h), \\ \overline{L}_f^m &= \frac{\sum_{\{x_t, y_t\}} \mathbb{1}[d_t^{EAT}(x_t) < 0.5] \ell_f(x_t, y_t)}{\sum_{\{x_t, y_t\}} \mathbb{1}[d_t^{EAT}(x_t) < 0.5]}, \quad (2) \\ \overline{L}_f^h &= \frac{\sum_{\{x_t, y_t\}} \mathbb{1}[d_t^{EAT}(x_t) \geq 0.5] \ell_f(x_t, y_t)}{\sum_{\{x_t, y_t\}} \mathbb{1}[d_t^{EAT}(x_t) \geq 0.5]}. \end{aligned}$$

Here, $\alpha > 0$ is a margin hyper-parameter. Given the fixed model, EAT carefully allocate the expert and the model annotations so that it ensures that the average loss of the model on the triage-to-human data \overline{L}_f^h is much higher than the triage-to-model data \overline{L}_f^m . Finally, the learning objective of EAT is given by $L_{EAT} = L_d + L_m$.

3.2.3 Bi-weighting mechanism

The AL-based mechanism focuses on triage-to-human data to improve the model’s predictive ability. However, it may still struggle with hard data beyond the model capacity. In contrast, our proposed EAT-based mechanism prioritizes triage-to-model data and reserves hard data for human annotation, which allows the model to handle relatively easy data. However, models trained on hard data may have weaker predictive ability. To find the optimal trade-off, the bi-weighting mechanism aims to ensure that the model is trained properly and can annotate relatively easy data effectively.

The bi-weighting mechanism aggregates the advantages of prioritizing both triage-to-human and

trriage-to-model data by introducing a simple yet efficient method. It scores x_t using $d_t^{bi}(x_t)$, which is the product of $d_t^{AL}(x_t) \in [0, 1]$ and $d_t^{EAT}(x_t) \in [0, 1]$, to adjust the importance of data. This ensures that informative or hard data are assigned to humans for annotation, and that the model is trained properly. Considering the sample efficiency of AL is more advantageous when the model is trained based on a few labeled data, we further control the importance of the AL score $d_t^{AL}(x_t)$ by $\tau \geq 0$, which is the percentage of annotated data in the total data. Given a function $\beta(\tau) = e^{\tau - T_0}$ parameterized by a hyper-parameter $T_0 \in [0, 1]$ and the total amount of dataset $|X|$, the weight of $x_t \in X$ is given by the following equation.

$$\begin{aligned} d_t^{bi}(x_t) &= \left(d_t^{AL}(x_t)\right)^{\eta(t)} * d_t^{EAT}(x_t) \\ \eta(t) &= \beta(t/|X|). \end{aligned} \quad (3)$$

As such, the bi-weighting mechanism highlights the importance of $d_t^{AL}(x_t)$ in the early stages of the annotation process (i.e., when $\tau < T_0$) and decreases it as the annotation process progresses. This allows for efficient use of the AL score. Notably, the bi-weighting mechanism is efficient because it allows for the fusion of AL and EAT in a post-hoc manner.

Optimization of SANT. We optimize the model annotator f_t , EAT d_t^{EAT} , and AL method d_t^{AL} independently. The loss function for AL, denoted as L_{AL} (if applicable), is included in the overall loss function \mathcal{L} of SANT, which is the sum of L_f , L_{EAT} , and L_{AL} . Note that all human-annotated data from past human-model interactions are utilized for optimization. For optimization details, please refer to Appendix F.

4 Evaluation

In this section, we conduct extensive experiments to assess the effectiveness of SANT. Given limited annotation budgets, we evaluate if SANT is more desirable to obtain high-quality annotations, compared to existing annotation methods (cf. Section 4.2). Furthermore, we comprehensively analyze the advantages of SANT and uncover the characteristics of triage-to-human and triage-to-model data in data annotation task (cf. Section 4.3).

4.1 Experiment Setup

Tasks & Datasets. To assess SANT, we perform three annotation tasks, namely sentiment annotation, knowledge graph completion, and tagging.

These tasks have been highlighted as crucial in various web applications (Yao et al., 2021; Marcos-Pablos and García-Peñalvo, 2020; Tovstogan et al., 2020; Lee et al., 2022). Each task corresponds to a different type of classification problem: binary classification, multi-class classification, and multi-label classification, respectively. The difficulty level of these tasks ranges from easy to hard. To simulate the annotation process, we utilize three commonly used datasets: IMDB for sentiment annotation, WN18RR for knowledge graph completion, and CiteULike for scientific paper tagging. More details about the datasets and tasks can be found in Appendix D.

Evaluation Metrics. We have selected appropriate metrics to evaluate the quality of annotations. For sentiment annotation and knowledge graph completion tasks, we use accuracy as the metric. Similarly, we use Hit Ratio ($HR@10$) for the multi-label tagging task, as suggested in Chen et al. (2020a). Our evaluation pays special attention to the annotation quality of model-annotated data, but also discusses the annotation quality of overall data (including expert-annotated and model-annotated data).

Comparative methods. For a fair comparison, we use the same model annotator for both SANT and semi-automatic annotation, with the only difference being the addition of a data triage module. We consider both classic and advanced active learning methods, which focus on triage-to-human data, and **Random** strategy that allocates data randomly.

- **MaxEntropy** uses entropy-based strategy as data triage to assign data with high uncertainty to the expert (Xiao et al., 2023; Li et al., 2023).
- **Calibrated MaxEntropy** aims to measure uncertainty more precisely. We equip *MaxEntropy* with the model calibration, implemented by the temperature scaling (Guo et al., 2017) with the scalar default parameter (i.e., 1.5).
- **Ent-gn** uses an advanced AL method (Wang et al., 2022) as the data triage module. It selects data that leads to a lower upper-bound of test loss, where the entropy is used to compute the gradient norm in an unsupervised way.
- **Exp-gn** uses labels from past interactions to compute an expected empirical loss (Wang et al., 2022) instead of using entropy in *Ent-gn*.
- **LLM-based automatic annotation** involves using ChatGPT as a strong baseline for automatic annotation. We test various prompts, including zero-shot, few-shot, zero-shot CoT, and few-shot




Figure 2: Annotation quality of SANT and ChatGPT-based automatic annotation. The X-axis means the proportion of annotation budgets. As the annotation tasks become increasingly difficult (from task a to task c), human experts are indispensable to achieving high-quality annotation, despite the efficiency of adopting LLMs as annotators.

CoT. Refer to Appendix G for more details.

To analyze SANT’s advantages and uncover its characteristics, we consider the following ablation baselines used in Section 4.3.

- **SANT w/o AL** prioritizes the triage-to-model data only. It assigns easy-to-predict data to the model and hard-to-predict data to the expert.
- **SANT w/o EAT** prioritizes the triage-to-human data only, which is essentially semi-automatic annotation with the AL enhancement. The off-the-shelf *Exp-gn* is used in this case.

Implementation details. Considering the human-model annotation efficiency, light-weighted annotation models are always preferred (Desmond et al., 2021a; Chen et al., 2020a; Hedderich et al., 2021). Regarding the semi-automatic annotation implementation, we use customized model annotators for different annotation tasks. Specifically, we use FastText for sentiment annotation (Hedderich et al., 2021), the distributional model for the knowledge graph completion task (Roller et al., 2014; Kober et al., 2021), and the item tagging model for tagging (Chen et al., 2020a). Regarding the implementation of SANT, MLP is used as the backbone of EAT following Chen et al. (2020a). The hyperparameter T_0 in the bi-weighting mechanism is set to be 0.2 for all experiments without tuning. To simulate human annotation behavior, we mask out the ground truth in the datasets. the ground truth is revealed when the expert is selected by the data triage module. Moreover, to simulate the limited annotation budgets setting, we tune the annotation budgets from 10% to 90% with a step size being 10%, where 10% budgets mean the human can only annotate 10% data. More implementation details are shown in Appendix E, and engineering-related suggestions are in Appendix H.

4.2 Annotation Quality Evaluation

This section presents a comparison of the annotation quality of SANT with LLM-based automatic methods and semi-automatic methods. The results are presented in Figure 2 and Table 1.

Despite the efficiency of adopting LLMs as annotators, human experts are indispensable to achieving high-quality annotation. As illustrated in Figure 2, while ChatGPT’s performance on simple annotation tasks (i.e., sentiment annotation) is comparable to SANT, it falls behind SANT significantly when it comes to complex annotation tasks (i.e., knowledge graph completion and tagging). In particular, ChatGPT struggles with the tagging task due to the large number of candidate tags (i.e., 500). It often generates non-existent tags. Compared to SANT, which uses a light-weighted model architecture and a certain amount of annotation budgets, we suggest that experts play an indispensable role in achieving high-quality annotations, especially for challenging annotation tasks. Therefore, significant effort should be invested in researching methods that make full use of limited budgets to achieve high-quality annotations, which is our motivation.

SANT secures higher-quality annotations compared to semi-automatic annotation. Table 1 shows the annotation quality of SANT and baselines on the model-annotated data, excluding the influence of expert annotation. Compared to random-based method, almost all triage-based methods could achieve performance improvements. Such an exception motivates the necessity of clear division of duty between the human and model. This is beneficial for maximizing the annotation quality and optimizing the utilization of the limited budgets. Compared to AL-based semi-automatic annotation, SANT secures higher quality. Moreover, SANT consistently outperforms *Exp-gn*, which shares the same model annotator and AL-based mechanism, across different annotation budgets. As the annotation budgets increase, the performance gain of

Task	Triage Method	Triage Types	10%	20%	30%	40%	50%	60%	70%	80%	90%
Sentiment Annotation (<i>binary classification</i>)	Random	Random	82.90	82.38	82.24	82.74	82.15	82.23	82.27	82.49	82.44
	Calibrated MaxEntropy	Classic AL	84.28	85.42	87.45	89.31	92.58	96.17	96.17	96.10	96.16
	MaxEntropy	Classic AL	83.15	85.09	87.03	89.61	93.21	95.94	95.93	95.91	95.93
	Ent-gn	Advanced AL	<u>84.37</u>	85.59	<u>87.53</u>	<u>90.76</u>	93.00	96.32	<u>96.33</u>	96.32	<u>96.47</u>
	Exp-gn	Advanced AL	84.38	<u>85.57</u>	87.47	<u>90.76</u>	92.98	<u>96.29</u>	<u>96.33</u>	<u>96.40</u>	96.43
	SANT (w/ Exp-gn)	Advanced AL + EAT	84.38	<u>85.57</u>	87.93	90.79	93.77	96.28	97.00	97.67	97.77
Knowledge Graph Completion (<i>multi-class classification</i>)	Random	Random	55.97	55.78	55.92	55.81	55.87	57.08	57.03	57.06	56.82
	Calibrated MaxEntropy	Classic AL	<u>56.66</u>	58.91	60.75	<u>64.17</u>	64.48	64.72	64.59	64.59	64.29
	MaxEntropy	Classic AL	56.88	<u>58.90</u>	<u>61.27</u>	62.95	63.08	63.92	64.13	63.18	62.89
	Ent-gn	Advanced AL	54.48	56.08	59.34	60.15	65.19	<u>68.71</u>	75.49	78.09	79.41
	Exp-gn	Advanced AL	55.14	56.81	60.28	61.51	<u>65.40</u>	68.07	<u>76.90</u>	<u>81.76</u>	<u>81.41</u>
	SANT (w/ Exp-gn)	Advanced AL + EAT	55.14	56.81	62.83	66.32	71.78	78.67	85.03	87.31	87.16
Tagging (<i>multi-label classification</i>)	Random	Random	48.78	49.45	50.28	51.75	<u>49.58</u>	50.48	49.86	50.11	55.06
	Calibrated MaxEntropy	Classic AL	48.78	48.13	48.06	48.20	48.27	48.99	49.55	49.28	49.37
	MaxEntropy	Classic AL	48.78	48.20	48.40	48.40	48.47	48.68	49.24	49.27	49.35
	Ent-gn	Advanced AL	48.78	49.88	<u>51.34</u>	<u>52.94</u>	49.52	<u>52.02</u>	54.80	56.43	<u>58.43</u>
	Exp-gn	Advanced AL	48.78	49.90	50.94	52.89	49.50	50.24	54.48	55.47	57.84
	SANT (w/ Exp-gn)	Advanced AL + EAT	48.78	49.90	52.24	54.35	55.25	56.60	61.13	68.24	68.53

Table 1: The annotation quality on model-annotated data (%). The best performance is marked in **bold** and second best is underlined. Overall, SANT enjoys higher-quality annotations. We analysis the cons of SANT on very limited data in Section 4.3. More experiments are provided in Appendix, Table 5

SANT becomes more significant due to better training of the EAT and AL components. These experiments highlight that solely improving model performance through triage-to-human data (AL) may not be sufficient for optimizing annotation quality under limited budgets. Determining which data should be assigned to the model is crucial for improving annotation quality.

4.3 Analysis and Characteristics of SANT

We aim to find the key predictor for the annotation prediction of SANT: triage-to-human data or triage-to-model data. The former improves the predictive ability of the model via prioritizing informative data selected by AL, while the latter carefully selects easy-to-predict data for model annotation and prioritizes hard ones for expert annotation via EAT. We consider two ablation baselines that separate the impact of AL and EAT. The results are shown in Figure 3 and Table 2.

Prioritizing triage-to-human data for active learning can enhance model predictions, but this advantage is not always the case. We first study the effect of triage-to-human data by measuring the predictive ability. However, although conducting experiments under the same annotation budgets, models are trained on different human-annotated data and tested on different unlabeled data due to different data triage preferences of AL and EAT. This hinders the measurement of predictive ability. To address this, controlled experiments are conducted to assess the predictive ability of each model on the same extra test data. The results

are illustrated in Figure 3. AL performs better in predicting sentiment annotations but loses its advantages in completing knowledge graph tasks. On the other hand, EAT, which implies the anti-curriculum learning (Braun et al., 2017), or hard example mining (Shrivastava et al., 2016), achieves surprisingly good results in knowledge graph completion. Although prioritizing hard instances has been shown effective in training models (Mindermann et al., 2022; Shrivastava et al., 2016), our experiment does not find statistically significant evidence of its effectiveness. Therefore, determining which instances should be triaged to humans for training a good model remains a challenge. Our results suggest that future studies should consider incorporating hard instances.

Importance of triage-to-human and triage-to-model data varies at different annotation stage, which is effectively balanced by SANT. We further take both the triage-to-human and triage-to-model data into account simultaneously. Table 2 shows the annotation quality of SANT and its ablations on the model-annotated data². When working with very few budgets (10%-20%), using AL to utilize triage-to-human data is the optimal choice as *SANT w/o EAT* achieves higher quality on two-thirds of the datasets compared to *SANT w/o AL*. On the contrary, when the budgets increase, EAT, which estimates model misclassification, tends to be a more stable and accurate method for human-model collaborative annotation compared to AL.

²The results of annotation quality on the overall data are available in Appendix B.




Figure 3: Model predictive ability evaluation on the same extra test dataset. While prioritizing triage-to-human data by AL has some advantages over triage-to-model data in promoting model predictive ability, it is not always the case (e.g., SANT w/o EAT loses its advantage in knowledge graph completion task).

Task	Annotation Method	Data Triage	10%	20%	30%	40%	50%	60%	70%	80%	90%
Sentiment Annotation	SANT w/o EAT	Advanced AL	84.38	85.57	87.47	90.76	92.98	96.29	96.33	96.40	96.43
	SANT w/o AL	EAT	<u>83.90</u>	<u>85.16</u>	<u>87.89</u>	<u>90.78</u>	<u>93.40</u>	96.13	<u>96.96</u>	<u>97.47</u>	<u>97.71</u>
	SANT	Advanced AL + EAT	84.38	85.57	87.93	90.79	93.77	<u>96.28</u>	97.00	97.67	97.77
Knowledge Graph Completion	SANT w/o EAT	Advanced AL	<u>55.14</u>	<u>56.81</u>	60.28	61.51	65.40	<u>68.07</u>	76.90	81.76	81.41
	SANT w/o AL	EAT	57.06	59.75	<u>62.75</u>	66.08	71.61	<u>78.57</u>	84.99	87.28	87.12
	SANT	Advanced AL + EAT	<u>55.14</u>	<u>56.81</u>	62.83	66.32	71.78	78.67	85.03	87.31	87.16
Tagging	SANT w/o EAT	Advanced AL	48.78	49.90	50.94	52.89	49.50	50.24	54.48	55.47	57.84
	SANT w/o AL	EAT	48.78	<u>49.81</u>	<u>51.14</u>	<u>52.97</u>	<u>55.23</u>	<u>56.53</u>	<u>61.08</u>	<u>68.20</u>	<u>68.50</u>
	SANT	Advanced AL + EAT	48.78	49.90	52.24	54.35	55.25	56.60	61.13	68.24	68.53

Table 2: The annotation quality of SANT and its ablation counterparts on the model-annotated data (%). Importance of triage-to-human (i.e., AL) and triage-to-model data (i.e., EAT) varies at different annotation stage. SANT effectively balances the advantages of the triage-to-human and triage-to-model data by integrating EAT and AL.

This shifts the focus to triage-to-model data. In conjunction with EAT, *SANT w/o AL* achieves the better results compared to *SANT w/o EAT*, particularly in the middle range of annotation budgets (40%-70%). Therefore, the importance of triage-to-human and triage-to-model data varies at different annotation stage. To bring together two types of data, SANT utilizes the bi-weighting mechanism. We found that the annotation quality of SANT is optimal in most cases, indicating that SANT effectively allocates triage-to-human data to train the model and enhance its predictive capability, while also allocating triage-to-model data to improve the model annotation quality by providing easy-to-predict data for the model.

The key factor influencing the quality of annotations in SANT lies in the careful selection of data to be assigned for model annotation. According to Figure 3, prioritizing triage-to-human data using AL or EAT does not lead to a significant improvement in the model’s predictive ability. While determining which instances should be triaged to human is crucial for model training, our experiment does not find statistically significant evidence of its advantages. It is possible that the potential values of data have not been fully utilized (Tang and Huang, 2019; Mindermann et al., 2022). Therefore, triage-to-human data is expected

to have a limited contribution to improving annotation quality in SANT. Additionally, our analysis of Table 2 shows that *SANT w/o AL* has a more significant advantage in model annotation data quality compared to *SANT w/o EAT*. Therefore, we argue that carefully selecting easy-to-predict data for model prediction is crucial in significantly improving overall annotation accuracy, even though training the model on manual labeled hard-to-predict data may not result in a stronger predictive ability. This finding highlights that deciding which data to use for model prediction is more important than training itself, particularly as budgets increase. In some cases, the improvement in annotation quality can compensate for the disadvantages of an under-trained model with lower predictive ability. To sum up, prioritizing triage-to-model data is a pivotal factor in determining the success of SANT.

5 Conclusion

We present the first work to study the roles of both triage-to-human and triage-to-model data. We formulate semi-automatic data annotation as a data triage problem via SANT, achieved by optimizing two novel mechanisms (i.e., EAT and bi-weighting mechanisms) in utilizing limited annotation budgets. Empirical experiments show that SANT sig-

nificantly outperforms the existing semi-automatic annotation methods. Although the effectiveness of LLM-based annotator, we find that experts are still indispensable to achieving high-quality annotations, especially on challenging data annotation tasks. Coordinately delegating the human and the model as a team can be a feasible solution. We also observe that relying on triage-to-human data to create a better model may not be the most effective approach for high-quality annotations. It may be more beneficial to identify triage-to-model data.

Limitations

Scale of Model Annotator. One limitation of our work is that we use lightweight models for annotation due to limited computational resources. While lightweight models are commonly employed in data annotation (Chen et al., 2020a; Desmond et al., 2021a), we do not investigate the effectiveness of larger language models as the model annotator in this study. However, SANT, based on lightweight models, still outperforms the LLM-based automatic annotator.

Model annotation cost. While the use of models for data annotation may incur computational costs, our current research focuses solely on human annotation costs, neglecting model annotation costs. This simplification, while not entirely reflective of real-world scenarios, is justified by the significantly higher cost of human annotation compared to model annotation (Li et al., 2023; Gilardi et al., 2023). Additionally, the light-weighted models employed in this study further minimize the model annotation cost, rendering it negligible. To enhance the practicality of our work, future research will incorporate the cost of model annotation.

Human annotation cost simulation. Human annotation cost on each piece of data may not necessarily be constant. The cognitive abilities of different humans and the difficulty level of different data can both impact the human annotation cost. Our paper, like the previous research, assumes that the human annotation cost on each data is constant. In future research, constructing a user simulator to simulate human cognitive abilities and evaluating the difficulty of the data could be the next research direction to improve our work.

Performance under extremely limited budgets. Our method may not exhibit a substantial advantage over the best baseline under extremely limited budgets. However, it consistently performs

at least on par, and often slightly better in this scenario. This, however, does not diminish the substantial effectiveness advantage demonstrated by our method in other scenarios. We have conducted a thorough analysis of the reasons behind our method’s limited performance improvement under extremely limited budgets. This analysis will guide our future work in developing targeted improvements to further enhance the effectiveness of our method.

Data annotation can be an intricate and complex engineering problem. While our work focuses on the data allocation problem and introduces the SANT framework for human-machine cooperative annotation, real-world deployment of SANT necessitates further systematic engineering considerations. These aspects, often overlooked in both our paper and other research on data allocation or data annotation, are addressed in Appendix H, where deployment suggestions are provided.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No. 62272330); in part by the Fundamental Research Funds for the Central Universities (No. YJ202219); in part by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant (No. MSS24C004).

References

- Jonathan F Bard. 2013. *Practical bilevel optimization: algorithms and applications*, volume 30. Springer Science & Business Media.
- Stefan Braun, Daniel Neil, and Shih-Chii Liu. 2017. A curriculum learning method for improved noise robustness in automatic speech recognition. In *2017 25th European Signal Processing Conference (EU-SIPCO)*, pages 548–552. IEEE.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Christopher Bryant, Mariano Felice, and Ted Briscoe. 2017. [Automatic annotation and evaluation of error types for grammatical error correction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 793–805, Vancouver, Canada. Association for Computational Linguistics.

- Jhonny Cerezo, Felipe Bravo-Marquez, and Alexandre Henri Bergel. 2021. [Tools impact on the quality of annotations for chat untangling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 215–220, Online. Association for Computational Linguistics.
- Nicolo Cesa-Bianchi, Claudio Gentile, and Francesco Orabona. 2009. Robust bounds for classification via selective sampling. In *Proceedings of the 26th annual international conference on machine learning*, pages 121–128.
- Nicolo Cesa-Bianchi, Claudio Gentile, Luca Zaniboni, and Manfred Warmuth. 2006. Worst-case analysis of selective sampling for linear classification. *Journal of Machine Learning Research*, 7(7).
- Jia Chen, Yike Wu, Shiwan Zhao, and Qin Jin. 2021. [Language resource efficient learning for captioning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1887–1895, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Xu Chen, Changying Du, Xiuqiang He, and Jun Wang. 2020a. Jit2r: A joint framework for item tagging and tag-based recommendation. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 1681–1684.
- Yilu Chen, Xiaojun Zeng, Xing Chen, and Wenzhong Guo. 2020b. A survey on automatic image annotation. *Applied Intelligence*, 50(10):3412–3428.
- Bhavana Dalvi, Sumithra Bhakthavatsalam, Chris Clark, Peter Clark, Oren Etzioni, Anthony Fader, and Dirk Groeneveld. 2016. Ike-an interactive tool for knowledge extraction. In *Proceedings of the 5th Workshop on Automated Knowledge Base Construction*, pages 12–17.
- Stefan Depeweg, Jose-Miguel Hernandez-Lobato, Finale Doshi-Velez, and Steffen Udluft. 2018. Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning. In *International Conference on Machine Learning*, pages 1184–1193. PMLR.
- Michael Desmond, Michelle Brachman, Evelyn Duesterwald, Casey Dugan, Narendra Nath Joshi, Qian Pan, and Carolina Spina. 2022. Ai assisted data labeling with interactive auto label. In *AAAI Conference on Artificial Intelligence*.
- Michael Desmond, Evelyn Duesterwald, Kristina Brim-ijoin, Michelle Brachman, and Qian Pan. 2021a. [Semi-automated data labeling](#). In *Proceedings of the NeurIPS 2020 Competition and Demonstration Track*, volume 133 of *Proceedings of Machine Learning Research*, pages 156–169. PMLR.
- Michael Desmond, Evelyn Duesterwald, Kristina Brim-ijoin, Michelle Brachman, and Qian Pan. 2021b. Semi-automated data labeling. In *NeurIPS 2020 Competition and Demonstration Track*, pages 156–169. PMLR.
- Bosheng Ding, Chengwei Qin, Linlin Liu, Yew Ken Chia, Boyang Li, Shafiq Joty, and Lidong Bing. 2023. [Is GPT-3 a good data annotator?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11173–11195, Toronto, Canada. Association for Computational Linguistics.
- Bosheng Ding, Chengwei Qin, Ruochen Zhao, Tianze Luo, Xinze Li, Guizhen Chen, Wenhan Xia, Junjie Hu, Anh Tuan Luu, and Shafiq Joty. 2024. Data augmentation using llms: Data perspectives, learning paradigms and challenges. *arXiv preprint arXiv:2403.02990*.
- Liat Ein-Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. 2020. [Active Learning for BERT: An Empirical Study](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7949–7962, Online. Association for Computational Linguistics.
- Yansong Feng and Mirella Lapata. 2008. [Automatic image annotation using auxiliary text information](#). In *Proceedings of ACL-08: HLT*, pages 272–280, Columbus, Ohio. Association for Computational Linguistics.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. Chatgpt outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Xingwei He, Zhenghao Lin, Yeyun Gong, A-Long Jin, Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan Duan, and Weizhu Chen. 2023. [Annollm: Making large language models to be better crowdsourced annotators](#).
- Michael A Hedderich, Lukas Lange, and Dietrich Klakow. 2021. Anea: distant supervision for low-resource named entity recognition. *the Practical Machine Learning For Developing Countries Workshop at ICLR*.
- Yan Hu, Iqra Ameer, Xu Zuo, Xueqing Peng, Yujia Zhou, Zehan Li, Yiming Li, Jianfu Li, Xiaoqian Jiang, and Hua Xu. 2023. Zero-shot clinical entity recognition using chatgpt. *arXiv preprint arXiv:2303.16416*.
- Chen Huang, Yiping Jin, Ilija Ilievski, Wenqiang Lei, and Jiancheng Lv. 2024a. [ARAIDA: Analogical reasoning-augmented interactive data annotation](#). In

- Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10660–10675, Bangkok, Thailand. Association for Computational Linguistics.
- Chen Huang, Peixin Qin, Wenqiang Lei, and Jiancheng Lv. 2023. [Reduce human labor on evaluating conversational information retrieval system: A human-machine collaboration approach](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10876–10891, Singapore. Association for Computational Linguistics.
- Chen Huang, Xinwei Yang, Yang Deng, Wenqiang Lei, JianCheng Lv, and Tat-Seng Chua. 2024b. [Co-matching: Towards human-machine collaborative legal case matching](#).
- Rebecca Hwa. 2000. Sample selection for statistical grammar induction. In *2000 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*, pages 45–52.
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Heinrich Jiang, Been Kim, Melody Guan, and Maya Gupta. 2018. To trust or not to trust a classifier. *NeurIPS*, 31.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431. Association for Computational Linguistics.
- Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. The INCEpTION platform: Machine-assisted and knowledge-oriented interactive annotation. In *COLING: System Demonstrations*, pages 5–9, Santa Fe, New Mexico. ACL.
- Thomas Kober, Julie Weeds, Lorenzo Bertolini, and David J. Weir. 2021. Data augmentation for hypernymy detection. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1034–1048.
- Oyku Deniz Kose and Yanning Shen. 2022. [Fairness-aware selective sampling on attributed graphs](#). In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5682–5686.
- Trausti Kristjansson, Aron Culotta, Paul Viola, and Andrew McCallum. 2004. Interactive information extraction with constrained conditional random fields. In *AAAI*, volume 4, pages 412–418.
- Stefan Larson, Adrian Cheung, Anish Mahendran, Kevin Leach, and Jonathan K Kummerfeld. 2020. Inconsistencies in crowdsourced slot-filling annotations: A typology and identification methods. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5035–5046.
- Yoonjoo Lee, John Joon Young Chung, Tae Soo Kim, Jean Y Song, and Juho Kim. 2022. Promptiverse: Scalable generation of scaffolding prompts through human-ai hybrid knowledge graph annotation. In *CHI Conference on Human Factors in Computing Systems*, pages 1–18.
- Aoxue Li, Weiran Huang, Xu Lan, Jiashi Feng, Zhen-guo Li, and Liwei Wang. 2020. Boosting few-shot learning with adaptive margin loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12576–12584.
- Minzhi Li, Taiwei Shi, Caleb Ziems, Min-Yen Kan, Nancy F. Chen, Zhengyuan Liu, and Diyi Yang. 2023. [Coannotating: Uncertainty-guided work allocation between human and large language models for data annotation](#).
- Yun Yao Li, Elmer Kim, Marc A Touchette, Ramiya Venkatachalam, and Hao Wang. 2015. Vinery: A visual ide for information extraction. *Proceedings of the VLDB Endowment*, 8(12):1948–1951.
- Aiwei Liu, Xuming Hu, Lijie Wen, and Philip S Yu. 2023. A comprehensive evaluation of chatgpt’s zero-shot text-to-sql capability. *arXiv preprint arXiv:2303.13547*.
- Christina Lohr, Johannes Kiesel, Stephanie Luther, Johannes Hellrich, Tobias Kolditz, Benno Stein, and Udo Hahn. 2019. [Continuous quality control and advanced text segment annotation with WAT-SL 2.0](#). In *Proceedings of the 13th Linguistic Annotation Workshop*, pages 215–219, Florence, Italy. ACL.
- Antonio Loquercio, Mattia Segu, and Davide Scaramuzza. 2020. A general framework for uncertainty estimation in deep learning. *IEEE Robotics and Automation Letters*, 5(2):3153–3160.
- Bo-Ru Lu, Nikita Haduong, Chia-Hsuan Lee, Zeqiu Wu, Hao Cheng, Paul Koester, Jean Utke, Tao Yu, Noah A. Smith, and Mari Ostendorf. 2024. [Does collaborative human-lm dialogue generation help information extraction from human dialogues?](#)
- David Madras, Toni Pitassi, and Richard Zemel. 2018. Predict responsibly: improving fairness and accuracy by learning to defer. *NeurIPS*, 31.
- Marian Marchal, Merel Scholman, and Vera Demberg. 2021. Semi-automatic discourse annotation in a low-resource language: Developing a connective lexicon for nigerian pidgin. In *Proceedings of the 2nd Workshop on Computational Approaches to Discourse*, pages 84–94.

- Samuel Marcos-Pablos and Francisco J García-Peñalvo. 2020. Information retrieval methodology for aiding scientific database search. *Soft Computing*, 24(8):5551–5560.
- Sören Mindermann, Jan M Brauner, Muhammed T Razzak, Mrinank Sharma, Andreas Kirsch, Winnie Xu, Benedikt Höltgen, Aidan N Gomez, Adrien Morisot, Sebastian Farquhar, et al. 2022. Prioritized training on points that are learnable, worth learning, and not yet learnt. In *International Conference on Machine Learning*, pages 15630–15649. PMLR.
- Hussein Mozannar and David Sontag. 2020. Consistent estimators for learning to defer to an expert. In *International Conference on Machine Learning*, pages 7076–7087. PMLR.
- Chenri Ni, Nontawat Charoenphakdee, Junya Honda, and Masashi Sugiyama. 2019. On the calibration of multiclass classification with rejection. *NeurIPS*, 32.
- Nastaran Okati, Abir De, and Manuel Rodriguez. 2021. Differentiable learning under triage. *NeurIPS*, 34.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. 2019. The algorithmic automation problem: Prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*.
- Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and X. Wang. 2020. A survey of deep active learning. *ArXiv*, abs/2009.00236.
- Stephen Roller, Katrin Erk, and Gemma Boleda. 2014. Inclusive yet selective: Supervised distributional hyponymy detection. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1025–1036.
- Claudia Schulz, Christian M Meyer, Jan Kiewewetter, Michael Sailer, Elisabeth Bauer, Martin R Fischer, Frank Fischer, and Iryna Gurevych. 2019. Analysis of automatic annotation suggestions for hard discourse-level tasks in expert domains. In *57th Annual Meeting of the Association for Computational Linguistics (Acl 2019)*, pages 2761–2772. Association for Computational Linguistics.
- Agam Shah and Sudheer Chava. 2023. Zero is not hero yet: Benchmarking zero-shot performance of llms for financial tasks. *arXiv preprint arXiv:2305.16633*.
- Amirreza Shirani, Franck Dernoncourt, Paul Asente, Nedim Lipka, Seokhwan Kim, Jose Echevarria, and Tamar Solorio. 2019. [Learning emphasis selection for written text in visual media from crowd-sourced label distributions](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1167–1172, Florence, Italy. Association for Computational Linguistics.
- Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. 2016. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 761–769.
- Hongjin Su, Jungo Kasai, Chen Henry Wu, Weijia Shi, Tianlu Wang, Jiayi Xin, Rui Zhang, Mari Ostendorf, Luke Zettlemoyer, Noah A. Smith, and Tao Yu. 2022. [Selective annotation makes language models better few-shot learners](#).
- Yiming Tan, Dehai Min, Yu Li, Wenbo Li, Nan Hu, Yongrui Chen, and Guilin Qi. 2023. Evaluation of chatgpt as a question answering system for answering complex questions. *arXiv preprint arXiv:2303.07992*.
- Ying-Peng Tang and Sheng-Jun Huang. 2019. Self-paced active learning: Query the right thing at the right time. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 5117–5124.
- Philip Tovstogan, Xavier Serra, and Dmitry Bogdanov. 2020. Web interface for exploration of latent and tag spaces in music auto-tagging. In *Proceedings of the 37th International Conference on Machine Learning, Vienna, Austria. ICML*.
- Hao Wang, Binyi Chen, and Wu-Jun Li. 2013. Collaborative topic regression with social regularization for tag recommendation. In *Twenty-Third International Joint Conference on Artificial Intelligence*.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023. Is chatgpt a good nlg evaluator? a preliminary study. *arXiv preprint arXiv:2303.04048*.
- Tianyang Wang, Xingjian Li, Pengkun Yang, Guosheng Hu, Xiangrui Zeng, Siyu Huang, Cheng-Zhong Xu, and Min Xu. 2022. Boosting active learning via improving test performance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8566–8574.
- Xin Wang and Siu Ming Yiu. 2021. Classification with rejection: scaling generative classifiers with supervised deep infomax. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 2980–2986.
- Yong Wang, Xiao-Ming Wu, Qimai Li, Jiatao Gu, Wangmeng Xiang, Lei Zhang, and Victor OK Li. 2018. Large margin few-shot learning. *arXiv preprint arXiv:1807.02872*.
- Zhuang Wang, Koby Crammer, and Slobodan Vucetic. 2012. Breaking the curse of kernelization: Budgeted stochastic gradient descent for large-scale svm training. *The Journal of Machine Learning Research*, 13(1):3103–3131.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Bryan Wilder, Eric Horvitz, and Ece Kamar. 2021. Learning to complement humans. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 1526–1533.
- Ruixuan Xiao, Yiwen Dong, Junbo Zhao, Runze Wu, Minmin Lin, Gang Chen, and Haobo Wang. 2023. [FreeAL: Towards human-free active learning in the era of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14520–14535, Singapore. Association for Computational Linguistics.
- Bingsheng Yao, Ishan Jindal, Lucian Popa, Yannis Katsis, Sayan Ghosh, Lihong He, Yuxuan Lu, Shashank Srivastava, James Hendler, and Dakuo Wang. 2023. [Beyond labels: Empowering human with natural language explanations through a novel active-learning architecture](#).
- Kaichun Yao, Chuan Qin, Hengshu Zhu, Chao Ma, Jingshuai Zhang, Yi Du, and Hui Xiong. 2021. An interactive neural network approach to keyphrase extraction in talent recruitment. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 2383–2393.
- Yangjun Zhang, Pengjie Ren, and Maarten de Rijke. 2021. A human-machine collaborative framework for evaluating malevolence in dialogues. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5612–5623.
- Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. 2003. [Learning with local and global consistency](#). In *Advances in Neural Information Processing Systems*, volume 16. MIT Press.

A More Related Work

Our research is also tied to data annotation and active learning. Furthermore, we address the data triage problem by utilizing the idea of learning to defer. To elaborate on these approaches, we provide a literature review and highlight the differences from existing methods.

Data Annotation. Computer-assisted annotation (or interactive annotation) becomes an alternative option. Such methods enable a model to facilitate human annotation by providing annotation suggestions on the fly (Klie et al., 2018; Lohr et al., 2019; Marchal et al., 2021). However, these

methods still require manual validation on each prediction, which is inefficient and time-consuming. Thus, it is not feasible for the setting with limited annotation budgets.

Active Learning. It targets to select informative or representative data that can maximize the performance of a model (Ren et al., 2020). One way to achieve this is through the uncertainty-based method (Desmond et al., 2021b), which selects data with high prediction uncertainty for human annotation. Density-based methods are also commonly used, but recent studies show that these two methods are highly correlated (Loquercio et al., 2020). Additionally, research on selective sampling tries to combine active learning and online learning (Cesa-Bianchi et al., 2006, 2009), where the model is allowed to adaptively query the label of an observed data on the fly, according to its potential to improve the model’s performance. Basically, selective sampling has shown its effectiveness in online large-scale kernel (Wang et al., 2012) and graph-based online learning (Kose and Shen, 2022). Semi-automatic annotation utilizes the online active learning (or selective sampling) methods to triage data for expert annotations and efficiently train a good model. However, our SANT goes beyond active learning by introducing the EAT mechanism, which identifies data that the model struggles to predict. SANT then employs the bi-weighting mechanism to combine EAT and active learning, enabling it to leverage both triage-to-human data (active learning) and triage-to-model data (EAT). This integration allows SANT to take full advantage of both methods, resulting in more efficient and accurate annotations.

Learning to Defer. The collaboration between humans and machines can be described as "learning to defer", where tasks are jointly completed by either the machine or deferred to humans for decision-making (Raghu et al., 2019; Okati et al., 2021; Madras et al., 2018). To achieve the overall goal, a triage strategy is employed to dynamically assign tasks to different collaborators, optimizing global utilization. Depending on the customized triage strategies, existing approaches assign data with high prediction uncertainty to humans (Ni et al., 2019), while out-of-distribution and noisy data to humans for decision-making (Wang and Yiu, 2021). Additionally, Jiang et al. (2018) assign data with unreliable model predictions to humans to mitigate prediction bias. Existing semi-automatic data annotation methods can also be incorporated

into the *learning to defer* framework, where the informative data are deferred to humans by AL methods. Recent work tries to assign data based on the uncertainty of LLM predictions (Li et al., 2023), where the parameters of LLM are frozen during the whole annotation process, as is the uncertainty measurement method. However, this approach hinders the model’s ability to improve its annotation results by learning from the data, relying heavily on LLM’s data annotation capability for general tasks. As highlighted by our work and other researchers (Xiao et al., 2023), the annotation outcomes of LLM in general tasks are not satisfying. To update the annotation model on the fly, we notice that applying "learning to defer" to data annotation tasks is particularly challenging. This is because existing works on learning to defer primarily focus on fully supervised settings where labels are available, whereas in the annotation task, labels are only accessible when a human is selected for annotation. This motivates us to develop a model that can learn and generalize from limited labeled data on the fly (cf. Section 3.2.3). Furthermore, we propose two tailored mechanisms, i.e., EAT and bi-weighting mechanisms, and experimentally reveal some insights for the data annotation task.

B More Evaluation Results of SANT on Overall Dataset

Tables 3 and 4 demonstrate the annotation quality of each method on the overall data. These findings align with those in Tables 1 and 2, indicating that proper allocation of data to both human and model workers results in higher-quality annotation. Moreover, SANT effectively balances the advantages of triage-to-human data and triage-to-model data by allocating triage-to-model data to improve model annotation quality through easy-to-predict instances.

Moreover, we report the more results of model annotation quality on two datasets, including SST-5 and FreeBase in Table 5. Due to financial limits, we compare our method with the advanced AL method (Exp-gn) and the random-based method. Based on the table results, it is clear that our method continues to be effective. Moreover, when compared to LLM-based annotators, our method can generate higher-quality annotations with some human assistance. This highlights the importance of human experts in achieving high annotation quality, even though LLMs may be efficient annotators. In

the revision, we will conduct more experiments to provide stronger validation.

C Ablation study on the proposed max-margin loss L_m

As shown in Section 3.2.2, EAT uses the max-margin loss L_m to take advantage of the partially labeled data. The loss is parameterized by a margin hyper-parameter τ to force the average loss on the assign-to-human data to be much higher than the assign-to-model data. In this section, we offer a detailed discussion of the loss itself and the sensitivity analysis. Specifically, the ablation study on L_m and the hyper-parameter sensitivity study on τ are given in Figure 4. Here, we take the sentiment annotation task as an example, and require *SANT w/o AL* to reject 50% of the data to the model in each batch, making sure our ablation studies are performed on the training data of the same size. We focus on the model performances on model-annotated data, marked as MA_{Accept} . Here, τ is tuned in $\{0, 0.1, 0.3, 0.5, 0.7, 1.0\}$. According to the results, we believe that the max-margin loss L_m helps train a better EAT. Moreover, the L_m is relatively less sensitive to its hyper-parameter τ . Thus, we simply choose $\tau = 0.3$.




Figure 4: Ablation study on L_m and hyper-parameter tuning analysis on τ . We simply choose $\tau = 0.3$ in our experiments.

D Details on Datasets

To comprehensively evaluate SANT, we use IMDB³, WN18RR⁴, and CiteULike⁵ to simulate sentiment annotation, knowledge graph completion, and tagging tasks, respectively. Those

³<https://www.kaggle.com/datasets/lakshmi25npathi/imdb-dataset-of-50k-movie-reviews>

⁴<https://paperswithcode.com/dataset/wn18rr>

⁵Following (Chen et al., 2020a; Wang et al., 2013), this data is used for tagging experiment. See <https://github.com/js05212/citeulike-a>

Task	Annotation Method	Data Triage	10%	20%	30%	40%	50%	60%	70%	80%	90%
Sentiment Annotation (binary classification)	Random	Random	84.61	85.90	87.57	89.64	91.08	92.89	94.68	96.50	98.24
	Calibrated MaxEntropy	Classic AL	85.85	88.34	91.22	93.59	96.29	98.47	98.85	99.22	99.62
	MaxEntropy	Classic AL	84.84	88.07	90.92	93.77	96.61	98.38	98.78	99.18	99.59
	Ent-gn	Advanced AL	<u>85.93</u>	88.47	<u>91.27</u>	<u>94.46</u>	96.50	98.53	<u>98.90</u>	99.26	<u>99.65</u>
	Exp-gn	Advanced AL	85.94	88.46	91.23	<u>94.46</u>	96.49	98.52	<u>98.90</u>	<u>99.28</u>	99.64
	SANT	Advanced AL + EAT	85.94	88.46	91.55	94.47	96.89	<u>98.51</u>	99.10	99.53	99.78
Knowledge Graph Completion (multi-class classification)	Random	Random	60.37	64.62	69.14	73.49	77.94	82.83	87.11	91.41	95.68
	Calibrated MaxEntropy	Classic AL	<u>60.99</u>	67.13	72.53	<u>78.50</u>	82.24	85.89	89.38	92.92	96.43
	MaxEntropy	Classic AL	61.19	<u>67.12</u>	<u>72.89</u>	77.77	81.54	85.57	89.24	92.64	96.29
	Ent-gn	Advanced AL	59.03	64.86	71.54	76.09	82.60	<u>87.48</u>	92.65	95.62	97.94
	Exp-gn	Advanced AL	59.63	65.45	72.20	76.91	<u>82.70</u>	<u>87.23</u>	<u>93.07</u>	<u>96.35</u>	<u>98.14</u>
	SANT	Advanced AL + EAT	59.63	65.45	73.98	79.79	85.89	91.47	95.51	97.46	98.72
Tagging (multi-label classification)	Random	Random	53.90	59.56	65.20	71.05	<u>74.79</u>	80.19	84.96	90.02	95.51
	Calibrated MaxEntropy	Classic AL	53.90	58.50	63.64	68.92	74.14	79.60	84.87	89.86	94.94
	MaxEntropy	Classic AL	53.90	58.56	63.88	69.04	74.24	79.47	84.77	89.85	94.94
	Ent-gn	Advanced AL	53.90	<u>59.90</u>	<u>65.94</u>	<u>71.76</u>	74.76	<u>80.81</u>	<u>86.44</u>	<u>91.29</u>	<u>95.84</u>
	Exp-gn	Advanced AL	53.90	59.92	65.66	71.73	74.75	80.10	86.34	91.09	95.78
	SANT	Advanced AL + EAT	53.90	59.92	66.57	72.61	77.63	82.64	88.34	93.65	96.85

Table 3: The annotation quality of SANT and baselines on overall data (%). We consider three annotation tasks with different difficulties. The best performance is marked in **bold** and second best is underlined.

Task	Annotation Method	Data Triage	10%	20%	30%	40%	50%	60%	70%	80%	90%
Sentiment Annotation	SANT w/o EAT	Advanced AL	85.94	88.46	91.23	94.46	96.49	98.52	98.90	99.28	99.64
	SANT w/o AL	EAT	<u>85.51</u>	<u>88.13</u>	<u>91.52</u>	94.47	<u>96.70</u>	98.45	<u>99.09</u>	<u>99.49</u>	<u>99.77</u>
	SANT	Advanced AL + EAT	85.94	88.46	91.55	94.47	96.89	<u>98.51</u>	99.10	99.53	99.78
Knowledge Graph Completion	SANT w/o EAT	Advanced AL	<u>59.63</u>	<u>65.45</u>	72.20	76.91	82.70	87.23	93.07	<u>96.35</u>	98.14
	SANT w/o AL	EAT	61.35	67.80	<u>73.93</u>	<u>79.65</u>	<u>85.81</u>	<u>91.43</u>	<u>95.50</u>	97.46	<u>98.71</u>
	SANT	Advanced AL + EAT	<u>59.63</u>	<u>65.45</u>	73.98	79.79	85.89	91.47	95.51	97.46	98.72
Tagging	SANT w/o EAT	Advanced AL	53.90	59.92	65.66	71.73	74.75	80.10	86.34	91.09	<u>95.78</u>
	SANT w/o AL	EAT	53.90	<u>59.85</u>	<u>65.80</u>	<u>71.78</u>	<u>77.62</u>	<u>82.61</u>	<u>88.32</u>	<u>93.64</u>	96.85
	SANT	Advanced AL + EAT	53.90	59.92	66.57	72.61	77.63	82.64	88.34	93.65	96.85

Table 4: The annotation quality of SANT and its ablation counterparts on the overall data (%).

three datasets are all English datasets and they formulate three different classification problems, including binary-class classification, multi-class classification, and multi-label classification problems. For the sentiment annotation, the annotator is required to label the binary sentiment based on the input movie review. As for the knowledge graph completion, the annotator is required to label the semantic relation of an input word pair. As for the tagging task, the annotator is required to label the tags of a scientific paper.

Specifically, IMDB is a common-used data with 50K movie reviews and binary sentiment classes. WN18RR we use is a subset of WN18 with no inverse relations. WN18RR contains 93003 triples, 40943 entities, and 11 relation types. Lastly, CiteULike (Chen et al., 2020a; Wang et al., 2013) is a commonly used tagging dataset, collected from CiteULike and Google Scholar. It includes 7288 authors’ 160272 citations on 8212 papers, together with 46391 different tags. Each paper in this dataset has its abstract, title, and multiple tags. The con-

catenation of them are used as the input in our experiments. Following (Chen et al., 2020a), we also use 500 most frequent tags in our tagging experiment.

Considering the human-model interaction efficiency, light-weighted annotation models or systems are always preferred (Desmond et al., 2021a; Chen et al., 2020a). In our experiments, we use the pre-trained GloVe⁶ (Pennington et al., 2014) for all tasks to encode the word representation without fine-tuning. For all the dataset, we first filter out words that are not contained in the pre-trained Glove embedding. For word pairs in WN18RR, we only retain those pairs that both words have Glove embeddings. Following previous works (Roller et al., 2014; Kober et al., 2021), the word pair embedding is simply obtained by concatenating the embedding of two words. As for the sentence embedding used in CiteULike and IMDB, we follow

⁶Our knowledge graph completion task is context-free. Considering the consistency of the experimental setup, we do not use contextual embedding like BERT embedding.

Task	Annotation Method	Data Triage	10%	20%	30%	40%	50%	60%	70%	80%	90%
Sentiment Annotation (<i>multi-class annotation</i>) SST-5	Zero-shot ChatGPT	None	41.30								
	Few-shot ChatGPT	None	53.85								
	Random	Random	30.02	33.59	36.64	37.56	38.23	38.79	39.33	39.54	39.95
	Exp-gn	Advanced AL	30.13	34.04	38.43	40.09	45.74	49.63	54.08	60.31	64.95
	SANT	Advanced AL + EAT	30.13	34.04	40.77	43.60	49.20	54.87	59.81	65.26	70.51
Knowledge graph completion (<i>multi-class annotation</i>) FreeBase	Zero-shot ChatGPT	None	38.85								
	Few-shot ChatGPT	None	51.60								
	Random	Random	22.46	35.19	42.77	48.54	52.50	55.31	57.75	60.26	61.88
	Exp-gn	Advanced AL	22.86	36.56	46.84	50.56	64.18	69.99	73.98	79.22	83.48
	SANT (Ours)	Advanced AL + EAT	22.86	36.56	45.03	52.17	58.79	74.53	77.04	83.77	88.13

Table 5: More results on the annotation quality of SANT and baselines on model-annotated data. Two more datasets are considered here, including the SST-5 and FreeBase. Due to financial constraint, we compare our method with the ChatGPT, advanced AL method (Exp-gn), and random-based method.

the FastText (Joulin et al., 2017) and obtain the sentence embedding by averaging the word embedding. We also remove the data if its 90% words (or more) have no Glove embeddings.

E Implementation details

E.1 Implementation of SANT

SANT is flexible to integrate fancy models. It shares the same model annotator with semi-automatic annotation. However, considering the efficiency requirement of the human-model interaction, light-weighted annotation models or systems are always preferred (Desmond et al., 2021a; Chen et al., 2020a). Referring to Chen et al. (2020a), MLP is used as the EAT. For Implementation, our EAT outputs a two-dimensional predictive distribution as the triage signal rather than a real number between 0 and 1. By this means, we avoid setting a threshold to map the number to binary values of 0 and 1. We also use the class weight to L_d to handle the potential class imbalance problem and penalize mistakes in the minority class. Specifically, the weight is calculated based on the cumulative class distribution of human annotations. As for the model annotator in SANT, we use the corresponding model used in the semi-automatic annotation method. All our codes are implemented in Pytorch with one Nvidia GeForce RTX 3090. the three-layer FC with ReLu activation and dropout is utilized. Moreover, we set $k = 3$ for retrieving top-k local information to integrate the input correlations. To ease the optimization, We also relax $\mathbb{1}[y_t \neq \max(f_t(x_t))]$ to $\mathbb{1}[HR@10 > 0]$ in L_d during the optimization of multi-label tagging annotator. Regarding the hyper-parameter T_0 in the bi-weighting mechanism, we set it to be 0.2 for all experiments without tuning.

Finally, to simulate the limited annotation budgets setting, we tune the annotation budgets from 10% to 90% with a step size of 10%, where 10% budgets mean the human can only annotate 10% data. We also require the human to annotate the first 1k data to initialize the model annotator for each annotation method. Note that when it comes to the Tagging task, the first 1k data and the 10% annotation budgets are almost equal. Therefore, in the case of 10% annotation budgets, all the budgets (i.e., the first 1k data) are used for human annotations. The corresponding results of the different triage-based methods in Table 1 and Table 2 are the same.

E.2 Implementation of Baselines

Following Wang et al. (2022), our AL-based methods update the task model on the fly and output the model at the end of interactive data selection/annotation. As for *MaxEntropy* and *Calibrated MaxEntropy*, we measure uncertainty in sentiment annotation and knowledge graph completion, entropy is used, following a recent annotation method (Desmond et al., 2021b; Li et al., 2023). In the case of the tagging task, which is a multi-label classification task, *total entropy* (Depeweg et al., 2018) is used. Next, a dynamic mechanism (Wilder et al., 2021) is used to assign the input data, rather than adjusting the threshold. Specifically, an data x_t is assigned to human if $(1 - d_t^{AL}(x_t)) \max(f_t(x_t)) < d_t^{AL}(x_t)$ holds. Here, $\max(\cdot)$ outputs the maximal value of the input vector. By this means, the human receives data with sufficiently high AL score, or sufficiently uncertain model output. As for *Exp-gn* and *Ent-gn*, the outputs of the active learning method (Wang et al., 2022) can not be easily normalized. In this case, half the data in a batch with the highest AL

scores⁷ are assigned to the human and the rest to the model.

E.3 Implementation of Human Simulation

We consider human-machine cooperative data annotation under limited budgets, specifically a budget that restricts the number of human annotators that can be employed to annotate the entire dataset. In this case, our experiments focus solely on estimating the cost of human annotation, as it significantly outweighs the cost of model annotation, as evidenced by previous studies (Li et al., 2023; Gilardi et al., 2023). In particular, following previous works on data annotation (Chen et al., 2020a; Desmond et al., 2021a; Hedderich et al., 2021) and human-model interaction (Hwa, 2000; Kristjansson et al., 2004), we assume a consistent cost for human annotation per data point. This allows us to estimate the overall annotation budget using the number of human-annotated data points. To simulate human annotation behavior, we mask out the ground truth in the datasets. The ground truth is revealed when the human is selected by the data triage module. Moreover, to simulate the constraints of limited annotation budgets, we adjust the budget allocation from 10% to 90% of the dataset, with increments of 10%. This means that under a 10% budget, human annotators can only label 10% of the total dataset. This range of budget constraints allows us to assess the performance of SANT under varying budget limitations.

F Optimization Details of SANT

In SANT, whenever the human reveals the ground truth of x_t , the model f_t is updated with L_f serving as the learning loss for the model (e.g., NLLLoss). However, our research focus is not on building a better model annotator. Therefore, we use different task-specified off-the-shelf model annotators in our experiments.

To simplify the optimization process, we optimize the model annotator f_t , EAT d_t^{EAT} , and AL method d_t^{AL} independently. The loss function for AL, denoted as L_{AL} (if applicable), is included in the overall loss function \mathcal{L} of SANT, which is the sum of L_f , L_{EAT} , and L_{AL} . It is important to note that all human-annotated data from past human-model interactions are utilized for optimization.

⁷Authors show the larger the amount of data queries, the better the AL effect. Considering the human annotation labors, we assign half data to be labeled by the human in experiments.

During optimization, three things are taken into consideration. First, the output of the data triage module is relaxed to continuous values over $[0, 1]$ to avoid the integer optimization problem. Second, the top-k retrieval and data triage modules are implemented using the Gumbel-softmax reparameterization trick (Jang et al., 2016) to ease the non-differentiable problem. Finally, the overall loss function is a bi-level optimization problem (Bard, 2013), where optimizing the data allocation (i.e., d_t^{AL} and d_t^{EAT}) is nested within another problem of optimizing the model annotator f_t . Thus, we update f_t and data allocation independently and iteratively in a coordinate-descent style. Specifically, let θ_f^k , $\theta_{d^{EAT}}^k$, and $\theta_{d^{AL}}^k$ be the network parameters for f , d^{EAT} , and d^{AL} at optimization iteration k . The update procedure is given in Eq.(4). In the k^{th} optimization iteration, we first optimize the parameter of model f with the parameter of d^{EAT} and d^{AL} fixed. In the $(k + 1)^{th}$ epoch, we exchange and optimize d^{EAT} and d^{AL} with f fixed, and so on.

$$\begin{aligned}\theta_f^{k+1} &= \theta_f^k - \nabla_f \mathcal{L} \left(f^k, (d^{EAT})^k, (d^{AL})^k \right), \\ \theta_{d^{EAT}}^{k+1} &= \theta_{d^{EAT}}^k - \nabla_{d^{EAT}} \mathcal{L} \left(f^{k+1}, (d^{EAT})^k, (d^{AL})^k \right), \\ \theta_{d^{AL}}^{k+1} &= \theta_{d^{AL}}^k - \nabla_{d^{AL}} \mathcal{L} \left(f^{k+1}, (d^{EAT})^k, (d^{AL})^k \right).\end{aligned}\tag{4}$$

In our approach, EAT is optimized towards measuring the predictive ability of f_{t+1} instead of f_t . It is important to note that if the AL method does not contain a learnable parameter, such as entropy-based AL, it can be omitted from the optimization process.

G Automatic annotation via ChatGPT

Since ChatGPT has demonstrated strong zero-shot, few-shot, and CoT capabilities in various domains and tasks (Liu et al., 2023; Hu et al., 2023; Shah and Chava, 2023), we use ChatGPT for automatic data annotation experiments. Specifically, we test various prompting methods for ChatGPT, including zero-shot, few-shot, zero-shot CoT, and few-shot CoT prompts. The prompts for three annotation tasks are provided in Table 7. Due to the cost of using ChatGPT, we sample a subset for each task and then deliver it to ChatGPT for data annotation. Specifically, for the sentiment annotation task, we randomly sample 1000 data for ChatGPT annotation. As for the knowledge graph completion task, the WN18RR data contains multiple relation types.

Res.	KGC			SA			Tagging		
	SANT (re-anno.)	SANT (re-train+re-anno.)	SANT	SANT (re-anno.)	SANT (re-train+re-anno.)	SANT	SANT (re-anno.)	SANT (re-train+re-anno.)	SANT
10%	61.93	-	59.63	86.05	-	85.94	53.90	-	53.90
20%	68.15	-	65.45	88.73	-	88.46	60.15	-	59.92
30%	74.67	-	73.98	92.09	-	91.55	66.32	-	66.57
40%	79.72	-	79.79	95.01	-	94.47	71.91	-	72.61
50%	86.04	-	85.89	97.32	-	96.89	78.10	-	77.63
60%	91.49	93.93	91.47	98.64	99.09	98.51	82.63	87.20	82.64
70%	95.60	95.61	95.51	99.15	99.24	99.10	88.33	89.31	88.34
80%	97.46	97.48	97.46	99.55	99.53	99.53	93.64	93.66	93.65
90%	98.71	98.72	98.72	99.80	99.76	99.78	96.85	96.91	96.85

Table 6: The overall annotation quality of SANT after the model annotator (re-train and) re-annotating the data that are previously model-annotated. Here, '-' means no human budgets left when data run out.

As a result, we choose 100 data at random from each class for ChatGPT annotation. Last, ChatGPT fails at the tagging task, because the dataset CiteU-Like contains 500 different tags in our experiments. Despite our several attempts, ChatGPT is unable to select one or more tags from the 500 available ones based on the semantics of the input words. Therefore, we cut down the candidates to 200.

In practice, as ChatGPT has been banned in some countries, we use Monica (refer to <https://monica.im/>) for annotations, which is a Chrome extension powered by ChatGPT API. As free users of Monica, we have a daily usage limit. Thus, we require Monica to annotate multiple instances at a time (precisely, 20 at a time).

H Practical suggestions for the interactive annotation system implementation

To implement SANT into an annotation system in the practical scenario, some tricks might be helpful to further reduce human annotation labor and improve the annotation quality of both the human and model.

H.1 Improving user experience

Considering generating an annotated dataset is usually time-consuming and expensive for the human, building an annotation tool with a user-friendly UI into the system plays an important role. A previous study on the tool design shows that if offering a better user experience during the annotation (Cerezo et al., 2021), the annotation quality improves.

H.2 Providing annotation suggestions to human

When the human is selected to annotate data, the annotation system could attach the model predictions as suggestions to the human (Desmond et al.,

2022; Klie et al., 2018), avoiding the need to annotate data from scratch and thus reducing human cognitive labor.

H.3 On error-prone human annotations

Human annotations could be error-prone and inconsistent. It would usually hamper the overall annotation quality. Some works from other research domains are devoted to examining the error-prone human annotations, such as the crowd-sourced data annotation (Shirani et al., 2019; Larson et al., 2020). Considering it is not the research focus of this paper, we suggest the reader to those works before implementing a real human-model interactive system.

H.4 Post hoc correction for model annotation

The system might also involve more strategies to improve model annotations in a post hoc way. The system could use the model annotator from the final round, and re-annotate data, that was originally model-annotated at the early annotation rounds. If the final model in SANT is properly trained, the quality of the annotations might be further enhanced. If there are some budgets left for the human to annotate, the final model could be re-trained before the re-annotating, together with the newly human-annotated data.

Taking SANT for example, Table 6 demonstrates the improvement in annotation quality after the model annotator re-annotated the data that were previously model-annotated. The re-annotation does help to improve the quality, and such improvement could be further enhanced by the re-training procedure. However, model re-training and model re-annotating require extra time cost. In the real-world scenario, one should decide whether to use the two extra modules (i.e., the model re-training and

model re-annotating) based on the product form of his or her annotation system.

Task	Zero-shot	Zero-shot CoT
SA	<p>### Read the sentences and identify their sentiment. There are two options: negative or positive ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>	<p>### Read the sentences and identify their sentiment. You should think step by step and output the sentiment of each sentence. There are two optional sentiments: negative or positive ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>
KGC	<p>### Read the word pairs and identify their semantic relation. There are 11 options: entailment, hyponym, hypernym, member, synonym, antonym, ... ### 1. [INPUT_WORD1], [INPUT_WORD2] 2. [INPUT_WORD3], [INPUT_WORD4]</p>	<p>### Read the word pairs and identify their semantic relation. You should think step by step and output the semantic relation of each sentence. You have the following options: entailment, hyponym, hypernym, member, synonym, antonym, ... ### 1. [INPUT_WORD1], [INPUT_WORD2] 2. [INPUT_WORD3], [INPUT_WORD4]</p>
Tagging	<p>### Read the following sentences, choose one or more tags from the TagSet that correspond to each sentence's semantics. TagSet={...} ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>	<p>### Read the following sentences, choose one or more tags from the TagSet that correspond to each sentence's semantics. You should think step by step and output the tag list of each sentence. TagSet={...} ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>
Task	Few-shot	Few-shot CoT
SA	<p>### Read the sentences and identify their sentiment. There are two options: negative or positive Here are two examples that you can refer to. Example 1: Input sentence: What a script, what a story, what a mess! Sentiment: negative Example 2: Input sentence: This is a great movie. Too bad it is not available on home video. Sentiment: positive ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>	<p>### Read the sentences and identify their sentiment. You should think step by step and output the sentiment of each sentence. There are two optional sentiments: negative or positive Example 1: Input sentence: What a script, what a story, what a mess! Think step by step: He/she thinks the script and story are messy. Therefore, the sentiment is negative. Sentiment: negative Example 2: Input sentence: This is a great movie. Too bad it is not available on home video. Think step by step: He/she thinks the movie is great and he/she would like to watch it on home video. He/she really love this movie. Therefore, the sentiment is positive Sentiment: positive ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>
KGC	<p>### Read the word pairs and identify their semantic relation. You have the following options: entailment, hyponym, hypernym, ... Here are some examples that you can refer to. Example 1: Input word pair: 3D, film semantic relation: hyponym Example 2: Input word pair: ability, unfitnes semantic relation: antonym Example 3: Input word pair: abominably,atrociously semantic relation: synonym Example 4: Input word pair: acclaim, sanction semantic relation: entailment ... ### 1. [INPUT_WORD1], [INPUT_WORD2] 2. [INPUT_WORD3], [INPUT_WORD4] ...</p>	<p>### Read the word pairs and identify their semantic relation. You should think step by step and output the semantic relation of each sentence. You have the following options: hyponym, antonym, synonym, hypernym, entailment, member,... Here are some examples that you can refer to. Example 1: Input word pair: 3D, film Think step by step: 3D filming are a type of film, so the relationship between the two words is hyponym semantic relation: hyponym Example 2: Input word pair: ability, unfitnes Think step by step: In most cases, the term "unfitness" refers to anything or someone who is not appropriate for the task at hand, meaning that their ability is inadequate. Therefore, the relationship between the two words is antonym semantic relation: antonym Example 3: Input word pair: abominably,atrociously Think step by step: Abhorrently and atrociously both refer to something that is extremely unpleasant or unfair. Therefore, the relationship between the two words is synonym semantic relation: synonym Example 4: Input word pair: acclaim, sanction Think step by step: Acclaim is used to describe enthusiastic public praise. The verb sanction can be used to signify to grant permission or approval for (an activity). So the relationship between the two words is entailment semantic relation: entailment ... ### 1. [INPUT_WORD1], [INPUT_WORD2] 2. [INPUT_WORD3], [INPUT_WORD4]</p>
Tagging	<p>### Read the following sentences, choose one or more tags from the TagSet that correspond to each sentence's semantics. Here are one example that you can refer to. Example 1: Input sentence: the metabolic world of escherichia coli ... Tag list: ... TagSet={...} ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>	<p>### Read the following sentences, choose one or more tags from the TagSet that correspond to each sentence's semantics. You should think step by step and output the tag list of each sentence. Here are one example that you can refer to. Example 1: Input sentence: the metabolic world of escherichia coli is not small to elucidate ... Think step by step: ... Tag list: ... TagSet={...} ### 1. [INPUT_SENTENCE1] 2. [INPUT_SENTENCE2]</p>

Table 7: ChatGPT prompts for three annotation tasks. We consider four kinds prompt, including zero-shot, few-shot, zero-shot CoT, and few-shot CoT prompts.