

Cost-efficient Crowdsourcing for Span-based Sequence Labeling: Worker Selection and Data Augmentation

Anonymous ACL submission

Abstract

This paper introduces a novel crowdsourcing worker selection algorithm, enhancing annotation quality and reducing costs. Unlike previous studies targeting simpler tasks, this study contends with the complexities of label interdependencies in sequence labeling. The proposed algorithm utilizes a Combinatorial Multi-Armed Bandit (CMAB) approach for worker selection, and a cost-effective human feedback mechanism. The challenge of dealing with imbalanced and small-scale datasets, which hinders offline simulation of worker selection, is tackled using an innovative data augmentation method termed *shifting*, *expanding*, and *shrinking* (SES). Rigorous testing on CoNLL 2003 NER and Chinese OEI datasets showcased the algorithm’s efficiency, with an increase in F_1 score up to 100.04% of the expert-only baseline, alongside cost savings up to 65.97%. The paper also encompasses a dataset-independent test emulating annotation evaluation through a Bernoulli distribution, which still led to an impressive 97.56% F_1 score of the expert baseline and 59.88% cost savings. Furthermore, our approach can be seamlessly integrated into Reinforcement Learning from Human Feedback (RLHF) systems, offering a cost-effective solution for obtaining human feedback. This research not only addresses the challenges in worker selection for intricate NLP tasks but also paves the way for more economical and efficient human-in-the-loop systems.

1 Introduction

Crowdsourcing, the practice of obtaining labeled data from a multitude of contributors (Howe, 2006), has emerged as a pivotal tool in data collection for deep learning models. It offers a cost-effective alternative to expert labeling, making it especially valuable in today’s data-driven research landscape (Nowak and Ruger, 2010). While its application spans various domains, from image labeling to text classification (Venanzi et al., 2014), this paper

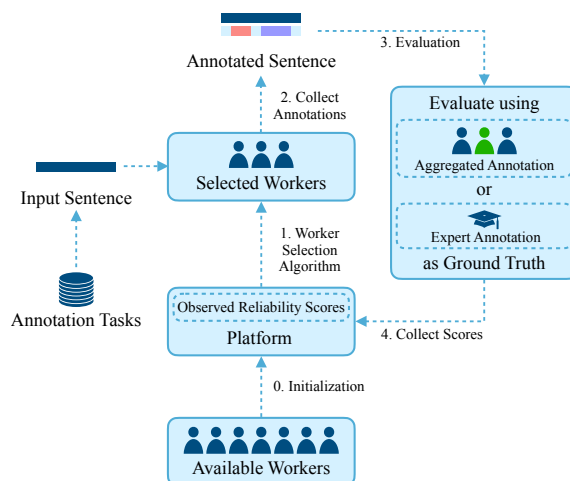


Figure 1: Our online worker selection framework for crowdsourcing.

narrows its focus on span-based sequence labeling tasks, which assign categorical labels to individual words within a sentence (Erdogan, 2010). Notable examples of such tasks include named entity recognition (NER) and opinion expression identification (OEI) (Collobert et al., 2011).

The inherent complexity of sequence labeling lies in the interdependencies of labels within a sequence. Unlike simpler tasks where labels are independent, sequence labeling requires contextual understanding, making it inherently more challenging (Rodrigues et al., 2014). Consequently, annotations from crowd workers, who might not possess the expertise of trained annotators, often exhibit reduced accuracy. This underscores the imperative to enhance annotation quality, a challenge that this study addresses.

A significant motivation driving this research is the potential application of a mixed feedback mechanism in Reinforcement Learning from Human Feedback (RLHF) systems. RLHF systems traditionally rely heavily on expert feedback, which, while accurate, is expensive and often not scalable (Casper et al., 2023). By integrating feedback

067 from both experts and aggregated crowd workers, 118
068 we can achieve a balance between accuracy and 119
069 cost. This hybrid approach not only maintains the 120
070 quality of feedback but also significantly reduces 121
071 the financial burden, making RLHF systems more 122
072 accessible and scalable. 123

073 Historically, research in this domain has con- 124
074 centrated on annotation aggregation (Rodrigues 125
075 et al., 2014; Nguyen et al., 2017; Simpson and 126
076 Gurevych, 2019), employing methods post data 127
077 collection. However, given the varied skill levels 128
078 among crowd workers, a proactive approach that 129
079 identifies and leverages the most accurate workers 130
080 during the data collection phase can significantly 131
081 enhance data quality. Termed as online worker se- 132
082 lection, this strategy involves iterative allocation 133
083 of a set budget across a pool of workers to opti- 134
084 mize annotation quality (Chen et al., 2013). This 135
085 dynamic process grapples with the uncertainty of 136
086 worker skill levels, necessitating a balance between 137
087 exploring new workers and exploiting currently 138
088 identified proficient ones.

089 In the context of sequence labeling, traditional 139
090 bandit-based algorithms (Rangi and Franceschetti, 140
091 2018) fall short due to the intricacies introduced 141
092 by label dependencies. These intricacies manifest 142
093 challenges in both annotation evaluation and aggre- 143
094 gation. To address the evaluation challenge, this 144
095 study employs the span-level F_1 score (Derczynski, 145
096 2016), a widely recognized metric, as the feedback 146
097 signal in the worker selection process. The core 147
098 **challenge** here is the accurate computation of the 148
099 F_1 score in the absence of expert annotations as a 149
100 reference. The objective is to minimize reliance on 150
101 costly expert annotations. For aggregation, while 151
102 the majority voting method is employed for its sim- 152
103 plicity and effectiveness, its reliability can be com- 153
104 promised when faced with divergent annotations 154
105 from different annotators.

106 The overarching **goal** of this research is to maxi- 155
107 mize the quality of annotations while minimizing 156
108 costs. This involves strategically replacing expert 157
109 ground truth labels with aggregated crowd-sourced 158
110 labels, ensuring that the overall F_1 score remains 159
111 high. Such replacements are made only when 160
112 there’s a high level of agreement among crowd 161
113 workers, indicating that expert evaluation might be 162
114 redundant for that particular sequence. The pro- 163
115 posed worker selection algorithm, as illustrated in 164
116 Figure 1, adopts an iterative approach: tasks are 165
117 assigned to a subset of workers, their annotations are

evaluated, and the resulting scores inform worker 118
selection in subsequent rounds. 119

120 However, real-world datasets present chal- 121
122 lenges due to their imbalanced nature and limited 123
124 scale (Rodrigues et al., 2014; Zhang et al., 2022). 125
126 Addressing these challenges, this paper introduces 127
128 a data augmentation method tailored for span-based 129
130 sequence labeling datasets. This method, designed 131
132 to emulate potential human annotation errors, en- 133
134 sures that aggregated annotations remain meaning- 134
135 ful. Three specific modifications, namely shift- 136
137 ing, expanding, and shrinking, are applied to ex- 138
139 pert annotations, generating a spectrum of potential 140
141 human annotations. This augmentation addresses 142
143 dataset limitations, facilitating the offline evalua- 144
145 tion of worker selection algorithms. 146

147 In summary, this paper’s contributions are mani- 148
149 fold¹: 149

- It presents the exploration of worker selection 136
for span-based sequence labeling tasks, recog- 137
nizing the unique challenges they present. 138
- It employs the span-level F_1 score, evaluated 139
by both experts and crowd workers, as a feed- 140
back mechanism, ensuring accurate worker 141
selection. 142
- It introduces a data augmentation technique 143
to counteract the limitations of real datasets, 144
enabling effective offline simulations. 145
- Through rigorous experimentation, it demon- 146
strates the efficacy of the proposed method, 147
achieving impressive F_1 scores while signifi- 148
cantly reducing expert annotation costs. 149

150 To provide a comprehensive understanding of 150
151 the background of this work, we briefly introduce 151
152 the related work in Appendix A. 152

153 2 Methodology

154 Consider an online crowdsourcing system that can 154
155 reach out to a group of crowd workers $W =$ 155
156 $\{w_1, w_2, \dots, w_N\}$. The workers are required to 156
157 provide sequential annotations to a set of sentences 157
158 $S = \{s_1, s_2, \dots, s_M\}$. More specifically, a worker 158
159 annotates a sentence by assigning a tag from a finite 159
160 possible tag set C (e.g., a set of BIO tags (Ramshaw 160
161 and Marcus, 1995)) to each word. An annotation 161

¹All resources, including source code and datasets, are made available for the broader research community at <https://anonymous.4open.science/r/CostEffi/>

on sentence s_i by worker w_j is a tag sequence $\mathbf{a}_{ij} = a_1 a_2 \dots a_k \dots a_l$ where $a_k \in C$ and l denotes the length of the sentence. We assume that every sentence is annotated by K different workers independently. We define a task as the process of annotating one entire sentence, and hence there are in total KM tasks. We seek to acquire an annotated dataset in which the average F_1 score of \mathbf{a}_{ij} is maximized. If we know which workers give the best annotations a priori, we can simply ask these workers to finish all the tasks. However, such information is unavailable in practice, and we aim to design an algorithm that learns the best workers throughout the crowdsourcing process.

In the beginning, we let each crowd worker annotate one sentence. We also ask the experts (e.g., well-trained linguists assumed to give the most precise annotations) to give one annotation for each of these sentences. Then we calculate the F_1 score of the annotation with the expert annotations as ground truth. We use these scores as the initial F_1 scores of workers. At each time step t after initialization (as illustrated in Figure 1), we select a subset of workers $W_t \subset W$ to do annotation, based on criteria discussed in Section 2.2. The size of the subset W_t should be neither too big nor too small (e.g., $0.3N$). We randomly choose a subset of sentences $S_t \subset S$, assign each $s_i \in S_t$ to K different workers in W_t , and collect their annotations $\mathbf{A}_i = \{\mathbf{a}_{i1}, \mathbf{a}_{i2}, \dots, \mathbf{a}_{iK}\}, \forall i \in \{1, 2, \dots, |S_t|\}$. To evaluate workers' F_1 scores on \mathbf{A}_i , one can use the expert annotations as the ground truth, which, however, can be very expensive (Iren and Bilgen, 2014). To cut down this cost, we reduce the usage of expert evaluations whenever crowd annotations are similar enough. We use the Fleiss' Kappa score κ to measure this similarity. The κ score ($\kappa \leq 1$) is a statistical measure of inter-annotator agreement. A larger value of κ indicates stronger agreement between the workers. κ score exceeding an empirical threshold indicates that the crowd workers reach a consensus on s_i . In that case, we aggregate \mathbf{A}_i with MV and use the aggregated annotation as the ground truth of sentence s_i . If the workers do not reach a consensus, we resort to expert annotations as ground truth. Next, we can calculate the F_1 scores of each $\mathbf{a}_{ij} \in \mathbf{A}_i$ and update the F_1 scores of the selected workers.

2.1 Problem Formulation

At time t , we obtain K crowd annotations \mathbf{A}_i on each sentence $s_i \in S_t$. We denote all annotations collected on S_t by $\mathcal{A}_t = \{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_{|S_t|}\}$. To simplify our expression, we use $F_1^{\text{Exp}}(\mathbf{a}_{ij})$ to represent the F_1 score of \mathbf{a}_{ij} using expert annotation as ground truth, and $F_1^{\text{MV}}(\mathbf{a}_{ij})$ to represent the F_1 score of \mathbf{a}_{ij} using the MV aggregation of $\mathbf{A}_i \in \mathcal{A}_t$ as ground truth. On collected annotation sets, $F_1^{\text{Exp}}(\mathbf{A}_i)$ denotes the average F_1 score of all $\mathbf{a}_{ij} \in \mathbf{A}_i$. Similarly, $F_1^{\text{Exp}}(\mathcal{A}_t)$ denotes the average F_1 score of all $\mathbf{A}_i \in \mathcal{A}_t$. As $F_1^{\text{Exp}}(\mathcal{A}_t)$ reflects the true accuracy of crowd annotations, our objective is to maximize the average expectation, or equivalently the cumulative expectation of $F_1^{\text{Exp}}(\mathcal{A}_t)$ over time T . We formulate this problem as a CMAB problem below:

$$\max \sum_{t=1}^T \mathbb{E}[F_1^{\text{Exp}}(\mathcal{A}_t)] \quad (1)$$

$$\text{s.t. } W_t \subset W, t \in \{1, 2, \dots, T\} \quad (2)$$

Since we have no information about workers' average F_1 scores, we need to balance exploring potentially better workers and exploiting the current best workers during worker selection. This tradeoff is extensively discussed in bandit literature where arms with unknown distributions form super-arms. The arms are associated with a set of random variables $X_{j,t}$ with bounded support on $[0, 1]$. Variable $X_{j,t}$ indicates the random outcome of arm j in time step t . The set of random variables $\{X_{j,t} | t \geq 1\}$ associated with arm j are independent and identically distributed according to certain unknown distribution D_j with unknown expectation $\bar{\mu}_j$. The platform plays a super-arm at each time step, and the reward of arms in it is revealed. These rewards are used as a metric for selecting the super-arm in future time steps. After enough time steps, the platform will be able to identify the best super-arm and keep playing it to maximize the overall reward. Similar to bandit terminologies, we call each worker $w_j \in W$ an arm and the worker subset $W_t \subset W$ a super-arm selected at t .

2.2 Worker Selection Algorithm

Specifically, there are three methods to calculate the reward of worker w_j at time step t as follows.

Expert Only This is a benchmark approach where the F_1 score is calculated using only expert

256 annotations as ground truth. This method provides
 257 intuitively the most accurate F_1 scores. The reward
 258 of worker w_j is defined as:

$$259 \quad \mu_j^{\text{Exp}}(t) = F_1^{\text{Exp}}(\mathbf{a}_{ij}(t)) \quad (3)$$

260 The expert-only method requires an expert annota-
 261 tion on every sentence, which is costly and usually
 262 not practical.

263 **Majority Voting (MV)** To reduce expert anno-
 264 tations, we aggregate \mathbf{A}_i for each sentence s_i , and
 265 use the aggregated annotation via MV as ground
 266 truth, i.e.,

$$267 \quad \mu_j^{\text{MV}}(t) = F_1^{\text{MV}}(\mathbf{a}_{ij}(t)) \quad (4)$$

268 **Expert+MV** When workers give very different
 269 annotations on the same sentence (usually when
 270 the task is difficult), one can be uncertain about
 271 the voted (and possibly noisy) ground truth. In this
 272 case, we want to resort to both crowd workers and
 273 experts. The choice is based on the well-known
 274 Fleiss’ Kappa score κ that can quantitatively evalu-
 275 ate the agreement of crowd workers. For each
 276 sentence s_i , if $\kappa(\mathbf{A}_i)$ is greater than a preset em-
 277 pirical threshold value τ , the reward of annotating
 278 workers is $F_1^{\text{MV}}(\mathbf{a}_{ij}(t))$. Otherwise, the reward is
 279 $F_1^{\text{Exp}}(\mathbf{a}_{ij}(t))$. In this way, MV is only used when
 280 the crowd workers can reach an agreement. Thus
 281 the reward is always calculated based on reliable
 282 ground truth. We summarize the reward of worker
 283 w_j as:

$$284 \quad \mu_j^{\text{Exp+MV}}(t) = \begin{cases} F_1^{\text{MV}}(\mathbf{a}_{ij}(t)), & \kappa(\mathbf{A}_i) > \tau \\ F_1^{\text{Exp}}(\mathbf{a}_{ij}(t)), & \kappa(\mathbf{A}_i) \leq \tau \end{cases} \quad (5)$$

285 The ϵ -Greedy, Thompson Sampling, and Com-
 286 binatorial Upper Confidence Bound (CUCB) are
 287 three effective algorithms to solve the CMAB
 288 problem. For each worker $w_j \in W$, both algo-
 289 rithms maintain a variable $\bar{\mu}_j(t)$ as the average
 290 reward (i.e., the average F_1 score) of worker w_j
 291 at time step t . CUCB additionally maintains a vari-
 292 able $T_j(t)$ as the total number of sentences worker
 293 w_j has annotated till time step t . Details of the
 294 worker selection algorithm with our **Exp.+MV** met-
 295 ric are shown in Algorithm 1. As for the selection
 296 criterion mentioned in the algorithm, ϵ -Greedy uti-
 297 lize a hyper-parameter ϵ which refers to the prob-
 298 ability of exploring random workers. Thus $1 - \epsilon$
 299 refers to the probability of exploiting the best work-

Measure	Chinese OEI	CoNLL 2003
# of Sentence	8047	4580
# of Worker	70	47
Span Length	5.05	1.51
Max	658	1626
Min	153	48
Range	505	1578
Mean	368	350
Median	332.5	230
SD	135.23	328.01
Variance	18286.52	107589.34
CV	36.71%	93.57%

Table 1: Statistics of the original datasets. Span lengths are averages. The terms SD and CV represent Standard Deviation and Coefficient of Variation respectively. The metrics Max, Min, Range, Mean, Median, SD, Variance, and CV pertain to the number of sentences annotated by each worker, indicating dataset imbalances.

ers till the current time step. Formally, W_t is se-
 lected with a random variable $p \in [0, 1]$ as below:

$$W_t = \begin{cases} \text{random } W_t \subset W, & p < \epsilon \\ \operatorname{argmax}_{W_t \subset W} \sum_{w_j \in W_t} \bar{\mu}_j, & p \geq \epsilon \end{cases} \quad (6)$$

Thompson Sampling samples from gaussian dis-
 tributions of workers’ rewards at each time step t ,
 and select workers which could maximize the total
 reward. CUCB handles the tradeoff by adding an
 item considering T_j and t to $\bar{\mu}_j$ like:

$$W_t = \operatorname{argmax}_{W_t \subset W} \sum_{w_j \in W_t} \left(\bar{\mu}_j + \sqrt{\frac{3 \ln t}{2T_j}} \right) \quad (7)$$

This makes workers with less annotations more
 likely to be selected as the algorithm proceeds. We
 provide a brief analysis in Appendix C. We explain
 on the application of our worker selection algo-
 rithms when building new datasets in Appendix F

2.3 Data Augmentation Method

We propose the data augmentation method to fa-
 cilitate the offline simulation of the crowdsourcing
 process, thus evaluating the worker selection algo-
 rithms. During offline simulation, when the worker
 selection strategy selects a certain worker to anno-
 tate a certain sentence, we can use the annotation
 in the original dataset if it exists. But if the selected
 worker did not annotate the sentence in the original
 dataset, we need to generate an annotation for the

	Shifting	Expanding	Shrinking
Expert	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm
Modified by 1 word	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm
Modified by 2 words	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm
Modified by 3 words	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm	今天的阳光是轻柔而温暖的 Today's sunshine is gentle and warm

Figure 2: An example of the three methods to generate annotations. Chinese characters and corresponding English words with red backgrounds indicate annotation spans. Transliteration of the Chinese sentence is *jīn tiān de yáng guāng shì qīng róu ér wēn nuǎn de*.

sentence. And the generated annotation should be in the same quality (depicted in F-score) as the real annotations by the worker. The generated annotation will be then used with the other annotations on the same sentence for majority voting.

Generating the missing annotations for each worker w_j is a great challenge when we expect the generated annotations to reflect the factual reliability of w_j . In other words, we expect the average F_1 score of each $w_j \in W$ to remain constant before and after augmenting the dataset with generated annotations. This is critical and difficult since real datasets are imbalanced and of small scale that cannot well support worker selection algorithms.

As there lack previous work on generating missing crowd annotations for span-based sequence labeling, we start with several naive algorithms such as randomly generating label sequences as annotations, and mixing expert annotations with completely incorrect (e.g., empty) annotations. But these algorithms either cannot produce annotations with expected F_1 scores, or generate confusing annotations which make later aggregation meaningless. This motivates us to design a data augmentation method specialized for span-based sequence labeling datasets.

Through our statistical analysis and observation on the real datasets, we characterized the 3 most common annotation error patterns. Due to space limitation, we defer the detailed analysis to Appendix D. Based on these analysis results, we propose a data augmentation method as follows: For each sentence $s_i \in S$, we modify the annotation span based on the expert annotation. We use three types of modifications to generate new annotation spans with different F_1 scores as illustrated in Figure 2. The goal of these modifications is to simulate varying annotation errors made by human annota-

tors.

Shifting We move both the left and the right border of the annotation span simultaneously in the same direction by one word per step.

Expanding We set one of the span borders fixed, and move the other border by one word per step to *increase* the length of the annotation span.

Shrinking We set one of the span borders fixed, and move the other border by one word per step to *decrease* the length of the annotation span.

We perform these modifications on a span multiple times, generating new annotation spans, until (1) the modified span does not overlap with the original one, (2) one of the span borders reaches an end of sentence or another span in the same sentence, or (3) the span length becomes 0.

For each sentence $s_i \in S$, s_i may contain multiple annotation spans. We perform modifications on each span in s_i , and find all combinations of spans to form possible sentence annotations. With these methods, we can imitate crowd annotations with different kinds of errors in practice. Next, for each worker $w_j \in W_{ti}$, if w_j has no annotation on s_i in the original dataset, we select one from all the expert and generated annotations on s_i .

We first calculate $\bar{\varphi}_j$ as the average F_1 score of all annotations by w_j on the original dataset, and then follow the detailed steps described in Algorithm 2 to do the selection. We aim to keep the overall F_1 score of w_j unchanged.

To better illustrate the procedure of the augmentation, we provide a running example in Appendix E.

Worker ID	Rnd. Gen. $ \Delta F_1 $	SES Only $ \Delta F_1 $	SES +Alg.2 $ \Delta F_1 $
25	2.83	6.69	0.01
52	8.15	10.83	0.00
46	3.83	13.48	0.00
43	10.02	11.21	0.00
18	9.87	12.84	0.00
50	16.69	10.71	0.00
12	47.18	10.52	0.00
Avg.	14.08	10.90	0.0014

Table 2: Comparisons between different data augmentation methods on the error of span-level exact F_1 score of every crowd worker. The error $|\Delta F_1|$ is calculated as the absolute difference between each worker’s F_1 score after augmentation and his real F_1 score. **Rnd. Gen.** is a naive augmentation method with random generated annotations. **SES Only** indicates the *shifting*, *shrinking*, and *expanding* method we proposed. **SES + Alg.2** means SES with Algorithm 2 which is our optimal method.

3 Experiments

3.1 Original Datasets

We compare our CMAB-based algorithms to several widely adopted baselines on two span-based sequence labeling datasets.


CoNLL 2003 The CoNLL 2003 English named-entity recognition dataset (Tjong Kim Sang and De Meulder, 2003) is a collection of news article from Reuters Corpus (Lewis et al., 2004). The dataset contains only expert annotations for four named entity categories (PER, LOC, ORG, MISC). Rodrigues et al. (2014) collected crowd annotations on 400 articles from the original dataset.

Chinese OEI The Chinese OEI dataset (Zhang et al., 2022) consists of sentences on the topic of COVID-19 collected from Sina Weibo², in which the task is to mark the spans of opinion expressions. The Chinese OEI dataset contains expert and crowd labels for two opinion expression categories (POS, NEG). Detailed statistics are shown in Table 1.


3.2 Data Augmentation

We augment both datasets with the method proposed in Section 2.3. According to Table 1, the most hard-working annotator in the OEI dataset

²<https://english.sina.com/weibo/>



(a) Chinese OEI



(b) CoNLL 2003

Figure 3: Cumulative regrets w.r.t time steps of all different worker selection methods.

provided annotations on 658 sentences, while the least one annotated only 153 sentences. On average, each crowd worker annotated 368 out of 8047 sentences in the Chinese OEI dataset. For the offline simulation of the worker selection process, we want every worker to annotate all 8047 sentences. Therefore we need to generate the missing $8047 - 368 = 7679$ annotations for every worker, on average. This also applies similarly to the CoNLL 2003 dataset.

Through our method, the average F_1 score of each $w \in W$ remains nearly unchanged before and after augmenting the original dataset with generated annotations³. Due to space limitation, we present comparisons of different augmentation algorithms with 10 sampled workers in Table 2. The complete results are deferred to Table 7 in the appendix. These results show that our **SES + Alg.2** method clearly outperforms the other baselines, producing almost the same F_1 scores for each worker as their original ones.

³The augmentation procedure takes about 2 hours on a computer with a 2.9 GHz Quad-Core Intel Core i7 CPU.

Method	Token-level			Span-level Exact			Span-level Prop.		
	P	R	F ₁	P	R	F ₁	P	R	F ₁
Oracle	65.69	83.99	70.00	78.15	72.23	74.96	87.97	80.03	83.82
Random	55.95	66.42	57.50	64.42	55.64	59.40	75.70	62.61	68.54
ϵ -G (Exp.)	64.94	80.48	68.56	75.24	68.16	71.34	85.85	76.79	81.06
ϵ -G (MV)	64.44	80.22	67.98	74.69	67.59	70.77	85.67	76.09	80.59
ϵ -G (Exp.+MV)	64.68	80.94	68.41	75.08	68.37	71.40	85.93	76.62	81.01
TS (Exp.)	64.94	79.88	68.51	75.64	68.31	71.57	85.02	75.71	80.09
TS (MV)	64.47	79.19	67.91	74.97	67.54	70.80	84.14	74.21	78.86
TS (Exp.+MV)	64.20	79.09	67.62	75.27	67.83	71.12	84.77	75.39	79.81
CUCB (Exp.)	65.65	80.34	69.24	75.94	69.12	72.20	86.17	77.22	81.45
CUCB (MV)	65.39	80.00	68.91	75.95	68.90	72.08	86.13	76.67	81.12
CUCB (Exp.+MV)	65.33	81.12	69.11	75.70	69.30	72.21	86.17	77.28	81.48

Table 3: Detailed P, R, and F₁ scores of all methods on the CoNLL 2003 dataset. All our algorithms perform significantly better than the Random (i.e., naive crowdsourcing) baseline.

3.3 Worker Selection

Baselines We test the **Exp.+MV** method with 4 baselines: **Oracle**, **Random**, **Exp.**, and **MV**. **Oracle** always selects the empirical best super-arm W^{opt} at every time step t . **Random** selects a different set of workers randomly at every time step t , which is equivalent to usual crowdsourcing procedure without worker selection. **Exp.**, **MV**, and **Exp.+MV** are CMAB-based algorithms introduced in Section 2.2. The CMAB-based algorithms are tested with CUCB, Thompson Sampling and ϵ -Greedy as the worker selection criterion respectively.

Regret as a Metric We evaluate our worker selection algorithms using cumulative regret, a metric indicating the performance deviation from the oracle’s selection defined as:

$$R(T) = \sum_{t=1}^T \left(\sum_{w_j \in W^{opt}} \bar{\mu}_j - \sum_{w_k \in W_t} \mu_k(t) \right) \quad (8)$$

In our experiments, we request 10 annotations per sentence, allowing CMAB-based algorithms to converge, and select 20 workers at each time step t . On the Chinese OEI dataset, setting the kappa threshold τ to 0.4 in **Exp.+MV** results in a 57.02% reduction in expert annotation cost, while a 0.65 threshold on the CoNLL 2003 dataset leads to a 43.83% cost reduction.


Results show **Random** consistently underperforms across datasets. On the Chinese OEI dataset,

Exp.+MV surpasses **MV**, albeit with higher regret than **Exp.**, justified by the substantial cost savings. On the CoNLL 2003 dataset, **Exp.+MV** even outperforms **Exp.**, suggesting crowd workers can provide valuable input for simpler tasks like NER. Overall, algorithms employing the CUCB criterion demonstrate superior performance, with **CUCB (Exp.+MV)** excelling in balancing cumulative regret and expert cost.

Effect of τ on F₁ and cost Next, we discuss how different kappa threshold values τ affect the average F₁ score of the produced annotation dataset. We test $\tau \in [0, 1]$ with a step of 0.05. In real datasets like CoNLL 2003 and Chinese OEI, the number of annotations per sentence is often quite small. To better fit the practical situations, we ask for 4 annotations on each sentence in the following experiments. Other settings remain unchanged. Since CUCB performs better than Thompson Sampling and ϵ -Greedy on both datasets, we display only the results from CUCB in later experiments.


On the Chinese OEI dataset, as illustrated in Figure 4, F₁ increases sharply with $\tau \in [0, 0.4]$. When $\tau = 0.4$, **Exp.+MV** achieves 99.47% F₁ score of **Exp.**, and saves 47.19% of the expert cost. The F₁ score goes up slowly until τ reaches 0.8. When $\tau = 0.8$, the F₁ score of **Exp.+MV** becomes exactly the same as the one of **Exp.**, and **Exp.+MV** still saves 6.6% of the expert cost.

The results on the CoNLL 2003 dataset are shown in Figure 5. Similarly, the F₁ score of



(a) F₁ score w.r.t τ (b) Expert usage w.r.t τ

Figure 4: F₁ scores of the produced annotations and usage of expert for annotation evaluations w.r.t the kappa threshold τ of the **Exp.+MV** method on the Chinese OEI dataset.



(a) F₁ score w.r.t τ (b) Expert usage w.r.t τ

Figure 5: F₁ scores of the produced annotations and usage of expert for annotation evaluations w.r.t the kappa threshold τ of the **Exp.+MV** method on the CoNLL 2003 dataset.

the produced annotation dataset grows fast as $\tau \in [0, 0.45]$. When $\tau = 0.45$, the **Exp.+MV** method already produce an annotation dataset with its F₁ reaching 99.86% of **Exp.**. At this point, **Exp.+MV** saves 88.57% of the expert cost. When $\tau = 0.65$, **Exp.+MV** outperforms **Exp.** with a 100.04% F₁ score and a 65.97% reduction in expert usage.

Our **CUCB (Exp.+MV)** worker selection algorithm eliminates the need for expert evaluation on every sentence. Instead, we harness crowd intelligence via our kappa-thresholded MV, producing datasets of comparable or even superior quality to those relying solely on expert evaluations.

Extended F₁ Metrics All of the F₁ scores in the previous experiments are span-level proportional scores calculated by the proportion of the overlap referring to the expert annotation (Zhang et al., 2022). To provide additional comparisons between different methods, we also invoke token-level and span-level exact P, R, F₁ scores as supporting metrics. We run the whole process from data augmentation to worker selection with all 3 metrics separately. The kappa threshold τ in **Exp.+MV** is set to 0.4 on the Chinese OEI dataset and 0.65 on

the CoNLL 2003 dataset. Detailed scores are listed in Table 3 and 6. The results show that **Exp.+MV** achieves scores as good as **Exp.** and much better than **MV**, which validates previous experiments and shows our worker selection methods are robust to different metrics.

Feedback Simulator We also test our worker selection methods with a feedback simulator. The simulator generates numerical feedback from *Bernoulli* distribution in annotation evaluations. This is to eliminate the varying level of difficulty in different tasks and evaluate our worker selection algorithms under more stable settings. Our algorithm achieves good results on the simulator. We put the definitions and results in Appendix B.

Effect on ML Models To further show the effect of our worker selection algorithm on the performance of machine learning models, we have run experiments with several widely-accepted models and provide the results in Table 5. We observe a consistent increment of F₁ score on the ML models, with our bandit-based worker selection algorithm. This validates that our worker selection algorithm may help improve the performance of ML models while saving budget on data crowdsourcing.

4 Conclusion

In this study, we introduced a CMAB-based worker selection strategy tailored for span-based sequence labeling tasks, leveraging the span-level F₁ with **Exp.+MV** as a feedback mechanism. To address the challenges posed by unbalanced and limited real datasets, we innovated a data augmentation method. This technique not only facilitates offline simulation but also mirrors the genuine annotation behaviors of workers closely.

Our empirical evaluations underscore the efficacy of the proposed method. On the Chinese OEI dataset, our approach achieved an impressive 99.47% F₁ score, translating to a substantial 47.19% reduction in expert costs. Similarly, on the CoNLL 2003 dataset, we observed a remarkable 100.04% F₁ score, with savings of up to 65.97% in expert costs, both benchmarks set against expert-evaluation-only baselines. Furthermore, our method demonstrated its robustness with a 94.86% F₁ score and a 65.97% reduction in expert costs on a data-free simulator. Our approach also boosts ML model performance, optimizing both accuracy and cost.

572 **Limitations**

573 In this paper, we provide theoretical analysis and
574 offline simulation results of our worker selection
575 algorithm. These results show that our algorithm
576 performs well. But due to the budget limitation,
577 we are unable to apply our algorithm on real online
578 crowdsourcing systems and test it with real-time
579 annotation tasks. Future work could explore ex-
580 tending our worker selection strategy to other do-
581 mains and tasks and in online environment, further
582 solidifying its role as a versatile and cost-effective
583 tool in the realm of crowdsourcing and beyond.

584 **References**

585 Arpita Biswas, Shweta Jain, Debmalya Mandal, and
586 Y. Narahari. 2015. [A truthful budget feasible multi-](#)
587 [armed bandit mechanism for crowdsourcing time critical](#)
588 [tasks](#). In *Proceedings of the 2015 International*
589 *Conference on Autonomous Agents and Multiagent*
590 *Systems*, pages 1101–1109.

591 Stephen Casper, Xander Davies, Claudia Shi,
592 Thomas Krendl Gilbert, Jérémy Scheurer, Javier
593 Rando, Rachel Freedman, Tomasz Korbak, David
594 Lindner, Pedro Freire, et al. 2023. [Open problems](#)
595 [and fundamental limitations of reinforcement](#)
596 [learning from human feedback](#). *arXiv preprint*
597 *arXiv:2307.15217*.

598 Wei Chen, Yajun Wang, and Yang Yuan. 2013. [Com-](#)
599 [binatorial multi-armed bandit: General framework](#)
600 [and applications](#). In *Proceedings of the 30th Inter-*
601 *national Conference on Machine Learning*, pages
602 151–159.

603 Ronan Collobert, Jason Weston, Léon Bottou, Michael
604 Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011.
605 [Natural language processing \(almost\) from scratch](#).
606 *Journal of Machine Learning Research*, 12(76):2493–
607 2537.

608 Leon Derczynski. 2016. [Complementarity, F-score, and](#)
609 [NLP evaluation](#). In *Proceedings of the Tenth Inter-*
610 *national Conference on Language Resources and*
611 *Evaluation*, pages 261–266.

612 Hakan Erdogan. 2010. [Sequence labeling: Generative](#)
613 [and discriminative approaches](#). In *Proceedings of the*
614 *9th International Conference on Machine Learning*
615 *and Applications*, pages 1–132.

616 Evrard Garcelon, Vashist Avadhanula, Alessandro
617 Lazaric, and Matteo Pirodda. 2022. [Top k ranking](#)
618 [for multi-armed bandit with noisy evaluations](#). In
619 *Proceedings of The 25th International Conference*
620 *on Artificial Intelligence and Statistics*, pages 6242–
621 6269.

622 Jeff Howe. 2006. [The rise of crowdsourcing](#). *Wired*
623 *magazine*, 14(6):1–4.

Tze Leung Lai and Herbert Robbins. 1985. [Asymptoti-](#)
624 [cally efficient adaptive allocation rules](#). *Advances in*
625 *Applied Mathematics*, 6(1):4–22. 626

David D. Lewis, Yiming Yang, Tony G. Rose, and Fan
Li. 2004. [RCV1: A new benchmark collection for](#)
627 [text categorization research](#). *Journal of Machine*
628 *Learning Research*, 5:361–397. 629 630

Feng Li, Jichao Zhao, Dongxiao Yu, Xiuzhen Cheng,
and Weifeng Lv. 2022. [Harnessing context for](#)
631 [budget-limited crowdsensing with massive uncertain](#)
632 [workers](#). *IEEE/ACM Transactions on Networking*,
633 30(5):2231–2245. 634 635

Rajeev Motwani and Prabhakar Raghavan. 1995. [Ran-](#)
636 [domized Algorithms](#). 637

Nikita Nangia, Saku Sugawara, Harsh Trivedi, Alex
Warstadt, Clara Vania, and Samuel R. Bowman. 2021.
[What ingredients make for an effective crowdsourc-](#)
638 [ing protocol for difficult NLU data collection tasks?](#)
639 In *Proceedings of the 59th Annual Meeting of the As-*
640 *sociation for Computational Linguistics and the 11th*
641 *International Joint Conference on Natural Language*
642 *Processing*, pages 1221–1235. 643 644 645

An Thanh Nguyen, Byron Wallace, Junyi Jessy Li, Ani
Nenkova, and Matthew Lease. 2017. [Aggregating](#)
646 [and predicting sequence labels from crowd annota-](#)
647 [tions](#). In *Proceedings of the 55th Annual Meeting of*
648 *the Association for Computational Linguistics*, pages
649 299–309. 650 651

Stefanie Nowak and Stefan Rüger. 2010. [How reliable](#)
652 [are annotations via crowdsourcing: A study about](#)
653 [inter-annotator agreement for multi-label image an-](#)
654 [notation](#). In *Proceedings of the International Con-*
655 *ference on Multimedia Information Retrieval*, page
656 557–566. 657

Lance Ramshaw and Mitch Marcus. 1995. [Text chunk-](#)
658 [ing using transformation-based learning](#). In *Third*
659 *Workshop on Very Large Corpora*. 660

Anshuka Rangi and Massimo Franceschetti. 2018.
[Multi-armed bandit algorithms for crowdsourcing](#)
661 [systems with online estimation of workers’ ability](#). In
662 *Proceedings of the 17th International Conference on*
663 *Autonomous Agents and MultiAgent Systems*, pages
664 1345–1352. 665 666

Filipe Rodrigues, Francisco Pereira, and Bernardete
Ribeiro. 2014. [Sequence labeling with multiple an-](#)
667 [notators](#). *Machine learning*, 95(2):165–181. 668 669

Filipe Rodrigues and Francisco C. Pereira. 2018. [Deep](#)
670 [learning from crowds](#). In *Proceedings of the Thirty-*
671 *Second AAAI Conference on Artificial Intelligence*
672 *and Thirtieth Innovative Applications of Artificial*
673 *Intelligence Conference and Eighth AAAI Symposium*
674 *on Educational Advances in Artificial Intelligence*,
675 pages 1611–1618. 676

677	Edwin Simpson and Iryna Gurevych. 2019. A Bayesian approach for sequence tagging with crowds . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing</i> , pages 1093–1104.	730
678		731
679		732
680		733
681		734
682		735
683	Yiwen Song and Haiming Jin. 2021. Minimizing entropy for crowdsourcing with combinatorial multi-armed bandit . In <i>IEEE Conference on Computer Communications</i> , pages 1–10.	736
684		737
685		738
686		739
687	Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition . In <i>Proceedings of the Seventh Conference on Natural Language Learning</i> , pages 142–147.	740
688		741
689		742
690		743
691		744
692	Long Tran-Thanh, Sebastian Stein, Alex Rogers, and Nicholas R. Jennings. 2014. Efficient crowdsourcing of unknown experts using bounded multi-armed bandits . <i>Artificial Intelligence</i> , 214:89–111.	745
693		746
694		747
695		748
696	Matteo Venanzi, John Guiver, Gabriella Kazai, Pushmeet Kohli, and Milad Shokouhi. 2014. Community-based bayesian aggregation models for crowdsourcing . In <i>Proceedings of the 23rd International Conference on World Wide Web</i> , page 155–164.	749
697		750
698		751
699		752
700		753
701	Xin Zhang, Guangwei Xu, Yueheng Sun, Meishan Zhang, Xiaobin Wang, and Min Zhang. 2022. Identifying Chinese opinion expressions with extremely-noisy crowdsourcing annotations . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics</i> , pages 2801–2813.	754
702		755
703		756
704		
705		
706		
707	Deniz İren and Semih Bilgen. 2014. Cost of quality in crowdsourcing . <i>Human Computation</i> , 1.	757
708		
709		
710	A Related Work	
711	Many studies (Rodrigues et al., 2014; Rodrigues and Pereira, 2018; Nangia et al., 2021) have used crowdsourcing for its efficiency and scalability. However, crowdsourcing suffers from the diversity of crowd workers’ expertise and effort levels that are hardly measurable to task requesters. Different approaches to improving the quality of collected data have been proposed and studied. For span-based sequence labeling tasks, prior studies mainly focus on annotation aggregation. Rodrigues et al. (2014) proposed CRF-MA, a CRF-based model with an assumption that only one worker is correct for any label. HMM-crowd from Nguyen et al. (2017) outperforms CRF-MA, but the effect of sequential dependencies is not taken into account. Simpson and Gurevych (2019) uses a fully Bayesian approach BSC which is proved to be more effective in handling noise in crowdsourced data. Aggregation methods are used <i>after</i> the data collection process completes. But we aim to assure data	758
712		759
713		760
714		761
715		762
716		763
717		764
718		765
719		766
720		767
721		768
722		769
723		770
724		771
725		772
726		773
727		774
728		775
729		776
		777
		778
		779
	quality and reduce cost <i>during</i> collecting. To this end, we focus on worker selection in our paper.	
	In online worker selection, we need to balance between exploring new workers and exploiting observed good workers. This exploration-exploitation tradeoff is extensively studied in the bandit literature (Lai and Robbins, 1985). In practice, we usually employ multiple crowd workers at the same time to finish the tasks more effectively. The combinatorial multi-armed bandit (CMAB) (Chen et al., 2013) models this circumstance. Biswas et al. (2015); Rangi and Franceschetti (2018) reformulate the problem as a bounded knapsack problem (BKP) and address it with the B-KUBE (Tran-Thanh et al., 2014) algorithm. Song and Jin (2021) introduce empirical entropy as the metric in CMAB and minimize the cumulative entropy with upper confidence bound (UCB) based algorithm. Li et al. (2022) consider the scalability of worker selection on large-scale crowdsourcing systems. These studies propose different methods under the CMAB settings, but on more complex span-based sequence labeling tasks there exists no discussion. We present the study of worker selection with CMAB on span-based sequence labeling tasks and show that our work performs well on the quality and efficiency of data collection.	
	B Feedback Simulator	
	The performance of crowd workers can vary across different types of annotation tasks. To evaluate the Exp.+MV worker selection method in more stable conditions without task-specific influence, we do not actually annotate the sentences, but directly use a worker’s average F_1 score to simulate his score on each sentence he annotates. The simulated scores are used as the numerical feedback for worker selection. Specifically, for each worker w , we calculate in advance two average F_1 scores for all of their annotations on the original dataset. The two F_1 scores for each worker are calculated using expert and majority vote (MV) evaluation respectively, denoted as $\bar{\varphi}_w^{Exp.}$ and $\bar{\varphi}_w^{MV}$. At each time step t , for every sentence s_i in the sentence set to be annotated S_t , we ask K different workers from the current selected workers W_t to annotate it. Then, we use a random value between 0 and 1 as the agreement level κ . If κ exceeds the threshold value τ (set to 0.4 in Exp.+MV), we independently generate feedback for the K workers from a Bernoulli distribution with a probability parameter set to $\bar{\varphi}_w^{MV}$. If	

Method	F ₁
Oracle	74.12
Random	65.12
Exp.	69.78
MV	66.80
Exp.+MV	68.29

Table 4: The overall span-level proportional F₁ scores of all methods with the feedback simulator.

Method	Original	w/ Our Alg.
LSTM-Crowd-cat	52.66	54.27
Bert-BiLSTM-CRF	52.14	54.51
Annotator-Adaptor	53.86	56.16

Table 5: Span-level exact F₁ scores of widely-accepted deep learning models on the Chinese OEI dataset. LSTM-Crowd-cat is from [Nguyen et al. \(2017\)](#). Bert-BiLSTM-CRF and Annotator-Adaptor are from [Zhang et al. \(2022\)](#). We provide results with and without our worker selection algorithm.

not, the feedback is generated from a Bernoulli distribution with a probability parameter set to $\varphi_w^{Exp.}$. The span-level average F₁ scores of the annotated dataset using different worker selection algorithm are shown in Table 4. Our feedback mechanism **Exp.+MV** for worker selection achieved comparable performance to the expert-only mechanism **Exp.** (68.29 versus 69.78), while in the same time reduced expert involvement in evaluation by 59.88% under the dataset-independent conditions.

C Regret Analysis

We provide a brief regret analysis of the worker selection framework assuming that we use the ϵ -greedy algorithm and that each worker’s reward follows a Bernoulli distribution.

The main proof follows the proof of Theorem 1 in ([Garcelon et al., 2022](#)). The key contribution here is that we need to specify that the evaluation signal (generated by majority voting) is a generalized linear model of workers’ true reward signal (generated by expert/oracle). To this end, we utilize the following form of the Chernoff bound which applies for any random variables with bounded support.

Lemma 1 (*Chernoff Bound* ([Motwani and Raghavan, 1995](#))) *Let X_1, X_2, \dots, X_N be independent random variables such that $x_l \leq X_i \leq x_h$ for*

Algorithm 1 The worker selection algorithm with the Expert+MV metric.

- 1: Let each worker $w_j \in W$ annotate a random sentence and initialize variable $\bar{\mu}_j$ with F₁ by expert evaluation
- 2: For each worker $w_j \in W$, initialize $T_j \leftarrow 1$
- 3: $t \leftarrow |W|$
- 4: **while** unannotated sentences exist **do**
- 5: $t \leftarrow t + 1$
- 6: Select $W_t \subset W$ based on certain criterion (e.g., (6), (7))
- 7: Split W_t into several disjoint subsets $\{W_{t1}, \dots, W_{ti}, \dots, W_{tn}\}$, each containing K workers
- 8: **for all** W_{ti} **do**
- 9: Let each $w_j \in W_{ti}$ annotate an sentence s_i and collect the annotations \mathbf{A}_i
- 10: **if** $\kappa(\mathbf{A}_i) > \tau$ **then**
- 11: Update T_j and $\bar{\mu}_j$ with $F_1^{MV}(\mathbf{a}_{ij}(t))$
- 12: **else**
- 13: Update T_i and $\bar{\mu}_j$ with $F_1^{Exp}(\mathbf{a}_{ij}(t))$
- 14: **end if**
- 15: **end for**
- 16: **end while**

all $i \in \{1, 2, \dots, N\}$. Let $X = \sum_{i=1}^N X_i$ and $\mu = \mathbb{E}(X)$. Given any $\delta > 0$, we have the following result:

$$P(X \leq (1 - \delta)\mu) \leq e^{-\frac{\delta^2 \mu^2}{N(x_h - x_l)^2}}. \quad (9)$$

For the purpose of our discussion, let $X_i \in \{0, 1\}$ be a binary random variable, where $X_i = 0$ denotes that worker i provides an incorrect solution, and $X_i = 1$ denotes that worker i generates a correct solution. Define $X = \sum_{i \in \mathcal{N}} X_i$.

We aim to approximate P_{MV} , which is the probability that the majority of the N workers provide the correct estimate. We apply the Chernoff Bound in Lemma 1 to P_{MV} . We can compute

$$\mathbb{E}(X) = \bar{p} = \frac{\sum_{i=1}^N p_i}{N}. \quad (10)$$

Based on (9), we let $\mu = \mathbb{E}(X)$, $\delta = \frac{N(\bar{p} - \frac{1}{2})}{\frac{N}{2} + N(\bar{p} - \frac{1}{2})}$,

Algorithm 2 The annotation selection algorithm.

- 1: For each worker $w_j \in W$, maintain (1) a variable $\hat{\varphi}_j$ as the average F_1 score of the selected annotations by w_j so far, (2) a set \mathbf{A}^j of selected annotations by w_j
 - 2: Generate all possible annotations \mathbf{A}_1^p on $s_1 \in S$, calculate $F_1^{\text{Exp}}(\mathbf{a}_{1k})$ for each $\mathbf{a}_{1k} \in \mathbf{A}_1^p$
 - 3: For each $w \in W$, initialize $\hat{\varphi}_j$ with the $F_1^{\text{Exp}}(\mathbf{a}_{1k})$ closest to $\bar{\varphi}_j$, and append the \mathbf{a}_{1k} to \mathbf{A}^j
 - 4: **for all** $s_i \in S \setminus s_1$ **do**
 - 5: Generate all possible annotations \mathbf{A}_i^p on $s_i \in S$, calculate $F_1^{\text{Exp}}(\mathbf{a}_{ik})$ for each $\mathbf{a}_{ik} \in \mathbf{A}_i^p$
 - 6: **for all** $w_j \in W$ **do**
 - 7: **if** $\hat{\varphi}_j > \bar{\varphi}_j$ **then**
 - 8: Update $\hat{\varphi}_j$ with the maximal $F_1^{\text{Exp}}(\mathbf{a}_{ik})$ less than $\bar{\varphi}_j$, and append \mathbf{a}_{ik} to \mathbf{A}^j
 - 9: **else**
 - 10: Update $\hat{\varphi}_j$ with the minimal $F_1^{\text{Exp}}(\mathbf{a}_{ik})$ greater than $\bar{\varphi}_j$, and append \mathbf{a}_{ik} to \mathbf{A}^j
 - 11: **end if**
 - 12: **end for**
 - 13: **end for**
-

$x_l = 0, x_h = 1$, and get the following result:

$$P_{\text{MV}} = P\left(X \geq \frac{N}{2}\right) = 1 - P\left(X \leq \frac{N}{2}\right) \geq 1 - e^{-\frac{\delta^2 \mu^2}{N}} \quad (11)$$

$$= 1 - e^{-\frac{N^2(\bar{p}-\frac{1}{2})^2}{[\frac{N}{2}+N(\bar{p}-\frac{1}{2})]^2} \frac{[\frac{N}{2}+N(\bar{p}-\frac{1}{2})]^2}{N}} \quad (12)$$

$$= 1 - e^{-\frac{N^2(\bar{p}-\frac{1}{2})^2}{N}} \quad (13)$$

$$= 1 - e^{-N\left(\frac{\sum_{i=1}^N p_i}{N} - \frac{1}{2}\right)^2} \quad (14)$$

Through approximating P_{MV} by its lower bound in (14), we can see that the evaluation signal (represented by P_{MV}) is an increasing function in each worker's capability p_i and twice-differentiable. That is, P_{MV} is a generalized linear function, which satisfies Assumption 3 in (Garcelon et al., 2022). Therefore, one can follow the proof of Theorem 1 in (Garcelon et al., 2022) that the ϵ -greedy algorithm yields a sub-linear regret with order $\tilde{O}(T^{2/3})$.

D Case Study of Annotation Errors

Based on our statistical analysis of the Chinese OEI dataset, we find that 74.80% of annotations have

different types of errors. And these annotation errors could be decomposed to three basic error types, namely Shifting, Expanding, and Shrinking (SES). In our data augmentation algorithm, we reversely used SES modifications and their combinations on the ground truth annotations to generate annotations with varying errors made by crowd workers. In this section, we provide a detailed characterization of human-made errors observed on annotated data with real cases to better motivate these modifications.

Shifting Some crowd annotation spans are as long as expert ones, but their positions are wrong. *Shifting* simulates this type of error. As depicted in Figure 6, both the expert span and the crowd span are three words long and of negative polarity. The difference is that the crowd span is shifted to the left by 2 words compared with the expert span. This type of error can be generated with *Shifting* on the expert annotations.

Expert	如果你感到有些沮丧或失落，你不妨试试运动。 If you feel slightly depressed or lost, you could try sports.
Crowd Worker	如果你感到有些沮丧或失落，你不妨试试运动。 If you feel slightly depressed or lost, you could try sports.

Figure 6: A case in which the crowd worker annotates a span with correct length and polarity but incorrect position.

Expanding *Expanding* is used to generate longer (than expert span) error spans. It might be intuitive that annotators barely make errors such as expanding to a very long span. However, in the case illustrated in Figure 7, the expert annotates five short spans separated by commas, while the crowd worker uses a very long span that covers the whole sentence, which is obviously not accurate. To simulate such human-made errors, we can expand an expert span to cover unnecessary words. Statistically, 4.03% of annotation errors are very long spans with more than 15 Chinese characters. So we do not set an upper bound of span length in *Expanding*.

Shrinking *Shrinking* is useful since crowd workers often ignore some words when annotating. As shown in Figure 8, the crowd worker failed to find all words expressing positive opinions.

Sometimes crowd workers ignore a whole span in expert annotations. This is why we set the lower

Expert	良好的身体素质是高效的保障，是成功的保障，是高水准生活的保障，是为社会多做贡献的保障，是生命的保障！ Good physical fitness is the guarantee of efficiency, of success, of a high standard of living, of contributing more to society, of life!
Crowd Worker	良好的身体素质是高效的保障，是成功的保障，是高水准生活的保障，是为社会多做贡献的保障，是生命的保障！ Good physical fitness is the guarantee of efficiency, of success, of a high standard of living, of contributing more to society, of life!

Figure 7: A case in which the crowd worker uses a very long span that covers the whole sentence.

Expert	我昨天在家里过得高效而又充实。 I had a productive and fruitful day yesterday at home.
Crowd Worker	我昨天在家里过得高效而又充实。 I had a productive and fruitful day yesterday at home.

Figure 8: A case in which the crowd worker does not annotate all words with polarity.

bound of span length to zero in *Shrinking*, which means we can shrink a span into no span.

These three types of errors may occur separately or combined in real crowd annotations. Such that an error could be both shifted and shrunk. This is why we use the combination of these three types of modifications to simulate human-made errors in our data augmentation algorithm.

E A Running Example of Data Augmentation

We here provide a running example to illustrate how an annotation for a certain worker on a certain sentence is generated with our proposed augmentation method. Suppose we have an English sentence:

Although he looked very depressed yesterday, he has already become much more cheerful now.

And an expert annotation:

Although he looked [NEGATIVE: very depressed] yesterday, he has already become [POSITIVE: much more cheerful] now.

If the crowd worker Sam has an annotation on this sentence in the original dataset, we use it directly in the augmented dataset. Otherwise, we generate an annotation for Sam with our data augmentation method.

When generating annotation for Sam, we follow the steps below:

- For each span in the expert annotation, we apply the Shifting, Expanding, and Shrinking (SES) modifications on it. After this step, we have several lists of annotation, each list contain annotations with only one modified span:
 - List 1, modifications of the first span, containing N_1 annotations:
 - Although he looked [NEGATIVE: very depressed] yesterday, he has already become much more cheerful now. # *Unmodified, span-level proportional F1 = 1.0*
 - Although he looked very [NEGATIVE: depressed yesterday], he has already become much more cheerful now. # *Shifting, F1 = 0.5*
 - Although he looked very depressed [NEGATIVE: yesterday ,] he has already become much more cheerful now. # *Shifting, F1 = 0*
 - ... # *Other Shifting modifications*
 - Although he [NEGATIVE: looked very depressed] yesterday, he has already become much more cheerful now. # *Expanding, F1 = 1.0*
 - ... # *Other Expanding modifications*
 - Although he looked very [NEGATIVE: depressed] yesterday, he has already become much more cheerful now. # *Shrinking, F1 = 0.5*
 - ... # *Other Shrinking modifications*
 - List 2, modifications of the second span, containing N_2 annotations:
 - Although he looked very depressed yesterday, he has already become [POSITIVE: much more cheerful] now. # *Unmodified, span-level proportional F1 = 1.0*
 - Although he looked very depressed yesterday, he has already become much [POSITIVE: more cheerful now]. # *Shifting, F1 = 0.6667*
 - Although he looked very depressed yesterday, he has already become much more [POSITIVE: cheerful now .] # *Shifting, F1 = 0.3334*
 - ... # *Other Shifting modifications*
 - Although he looked very depressed yesterday, he has already [POSITIVE:]

960 **TIVE: become much more cheer-**
 961 **ful]** now. # *Expanding, F1 = 1.0*
 962 ... # *Other Expanding modifications*
 963 – Although he looked very depressed
 964 yesterday, he has already become
 965 much **[POSITIVE: more cheerful]**
 966 now. # *Shrinking, F1 = 0.6667*
 967 ... # *Other Shrinking modifications*

968 2. We choose one annotation from each list, and
 969 combine them to generate an annotation with
 970 2 spans. This is done for all combinations
 971 of the annotations in the two lists. Note that
 972 if the two spans overlay with each other, we
 973 merge them into one span. After step 2, we
 974 have one list of annotations:

975 Combined List, containing less than
 976 or equal to $N_1 \times N_2$ annotations:
 977 • Although he looked **[NEG-**
 978 **ATIVE: very depressed]**
 979 yesterday, he has already
 980 become **[POSITIVE: much**
 981 **more cheerful]** now. #
 982 *span-level proportional F1 =*
 983 *1.0*
 984 • Although he looked very **[NEG-**
 985 **ATIVE: depressed yesterday]**
 986 , he has already become **[POS-**
 987 **ITIVE: much more cheerful]**
 988 now. # *span-level proportional*
 989 *F1 = 0.75*
 990 • Although he looked very
 991 depressed **[NEGATIVE:**
 992 **yesterday ,]** he has already
 993 become **[POSITIVE: much**
 994 **more cheerful]** now. #
 995 *span-level proportional F1 =*
 996 *0.5*
 997 ... # *Other combinations with F1*
 998 *ranging from 0 to 1.0*

999 3. We choose one annotation from the combined
 1000 list as Sam’s annotation on this sentence, ac-
 1001 cording to the following procedure:

- 1002 (a) Sam has an average F1 score $F_{ori} = 0.57$
 1003 on the original (real) dataset.
- 1004 (b) We have already got 10 annotations for
 1005 Sam in the augmented dataset, which has
 1006 an average F1 score $F_{aug_10} = 0.54$.
- 1007 (c) We are choosing annotation on the 11th
 1008 sentence for Sam.

- (d) We firstly select two annotations with the
 1009 closest F1 scores to F_{ori} from the com-
 1010 bined list, one higher than F_{ori} , and one
 1011 lower than F_{ori} , as candidate annotations.
 1012 In this case, the two annotations could
 1013 have F1 scores of 0.58 and 0.52 respec-
 1014 tively. 1015
- (e) If $F_{aug_10} > F_{ori}$, we choose the anno-
 1016 tation with the lower F1 score (0.52)
 1017 as Sam’s annotation on this sentence.
 1018 Otherwise, we choose the annotation
 1019 with the higher F1 score (0.58). This
 1020 is to ensure that the average F1 score
 1021 of Sam’s annotations in the whole aug-
 1022 mented dataset, F_{aug} , is as close to F_{ori}
 1023 as possible, which reflects Sam’s reliabil-
 1024 ity (i.e., performance). In this case, we
 1025 choose the annotation with F1 score of
 1026 0.58. 1027

1028 By generating the missing annotations in the
 1029 original dataset with the method above, we could
 1030 have an augmented dataset.

1031 F Explanation of Worker Selection on 1032 Building New Datasets

1033 When creating new datasets, we expect to have a
 1034 few (e.g. five) experts and a relatively large group
 1035 of (e.g. a hundred) crowd workers available for
 1036 annotation.

1037 At each time step, we select a group of (e.g. 20)
 1038 crowd workers, and request them to annotate a few
 1039 (e.g. 5) sentences, resulting in 4 crowd annotations
 1040 on each sentence. Now we calculate the agreement
 1041 of the annotations on each sentence, if the agree-
 1042 ment is high (e.g. greater than 0.4), we use the MV
 1043 aggregation of the crowd annotations as the ground
 1044 truth, and calculate the F1 scores of each worker’s
 1045 annotation. Otherwise, we ask an expert to give
 1046 an annotation on the sentence, and calculate the
 1047 F1 score of each worker on the expert annotation.
 1048 Note that the expert annotates only when the agree-
 1049 ment is low. After this time step, we have crowd
 1050 annotations on the sentences and their F-scores,
 1051 which can be used to update the average score of
 1052 each worker. This procedure is repeated until we
 1053 have enough annotations on every sentence.

1054 In other words, the Expert+MV approach works
 1055 with both crowd workers and experts available (e.g.
 1056 on an online system) when building datasets. The
 1057 Expert+MV is an iterative approach in which the
 1058 expert annotates when needed. And it saves the

Method	Token-level			Span-level Exact			Span-level Prop.		
	P	R	F ₁	P	R	F ₁	P	R	F ₁
Oracle	62.88	68.62	64.80	54.48	51.97	53.07	72.79	64.07	68.15
Random	58.49	57.30	57.42	43.99	35.50	39.18	69.01	52.36	59.55
ϵ -G (Exp.)	61.91	64.58	62.61	51.72	46.37	48.76	72.28	60.25	65.72
ϵ -G (MV)	60.87	63.52	61.55	48.72	44.66	46.37	70.15	58.94	64.05
ϵ -G (Exp.+MV)	61.76	64.46	62.47	49.14	45.35	46.96	71.21	59.92	65.08
TS (Exp.)	62.66	64.91	63.20	49.76	42.34	45.69	72.15	60.20	65.63
TS (MV)	59.82	61.90	60.25	44.81	40.71	42.36	67.72	56.05	61.34
TS (Exp.+MV)	61.66	64.03	62.23	47.20	42.36	44.49	70.66	59.07	64.35
CUCB (Exp.)	63.02	63.75	62.93	52.24	45.51	48.56	73.05	59.53	65.60
CUCB (MV)	61.94	62.09	61.55	49.57	44.39	46.66	71.22	57.59	63.68
CUCB (Exp.+MV)	62.83	63.62	62.75	51.31	45.60	48.16	72.48	59.33	65.25

Table 6: Detailed P, R, and F₁ scores of all methods on the Chinese OEI dataset. All our algorithms perform significantly better than the Random (i.e., naive crowdsourcing) baseline.

1059 cost of expert annotations by using the MV aggrega-
1060 tion of crowd annotations as the ground truth when
1061 possible. Our experiment results show that the Ex-
1062 pert+MV approach can save 47.19% of the cost of
1063 expert annotations on the Chinese OEI dataset, and
1064 65.97% on the CoNLL’03 dataset respectively.

1065 However, even in the case that no expert is avail-
1066 able, which means that Expert+MV falls back to
1067 MV, we can still observe that the MV approach out-
1068 performs the Random baseline (which is an equiva-
1069 lent of normal crowdsourcing procedure which as-
1070 signs an equal amount of sentences to each worker
1071 randomly) by a large gap. In this case, the MV
1072 approach saves 100% of expert annotation cost, but
1073 still produced crowd annotation with good qual-
1074 ity. Please refer to Table 3 and Table 6 for more
1075 detailed results.

Worker ID	Ori. F ₁	Rnd. Gen. F ₁	SES Only F ₁	SES +Alg.2 F ₁	Worker ID	Ori. F ₁	Rnd. Gen. F ₁	SES Only F ₁	SES +Alg.2 F ₁
25	62.90	60.07	69.59	62.89	37	37.15	96.10	26.79	37.16
32	60.87	41.37	68.79	60.87	13	36.19	31.62	25.14	36.20
42	53.88	4.37	66.57	53.88	20	36.11	71.44	25.02	36.12
5	52.07	50.74	60.76	52.06	64	35.97	65.66	25.39	35.97
55	50.70	30.24	61.13	50.70	63	35.22	75.40	24.73	35.22
2	50.53	91.99	60.92	50.53	6	35.15	65.74	25.00	35.16
52	50.08	41.93	60.91	50.08	10	34.63	51.28	25.08	34.64
17	49.82	43.73	35.82	49.82	66	33.75	60.98	24.99	33.75
57	49.25	13.17	35.59	49.25	53	32.90	27.51	24.78	32.89
11	49.04	53.71	35.19	49.03	4	32.72	8.40	24.77	32.72
26	48.89	5.17	35.59	48.82	21	32.19	73.47	24.78	32.19
36	48.71	15.53	35.27	48.70	62	32.16	48.71	24.89	32.16
46	48.67	44.84	35.19	48.67	1	32.10	34.42	24.96	32.10
29	48.60	95.39	35.21	48.60	41	31.94	77.55	24.88	31.93
35	47.07	23.64	35.34	47.07	51	31.78	68.07	24.85	31.78
49	46.80	60.30	35.27	46.80	31	31.61	29.44	24.59	31.61
54	45.63	18.74	34.45	45.64	8	31.05	28.55	24.76	31.05
14	45.13	60.99	34.54	45.13	67	30.91	95.51	24.22	30.91
43	44.93	34.91	33.72	44.93	58	30.70	21.64	23.96	30.70
7	44.37	23.89	33.50	44.37	65	30.61	4.51	24.17	30.60
59	44.36	72.37	33.61	44.37	38	30.47	4.82	24.11	30.47
23	43.38	4.85	33.58	43.38	28	29.86	2.63	24.00	29.86
56	43.37	41.96	33.31	43.37	45	29.38	36.13	24.15	29.38
0	41.60	66.81	28.19	41.61	30	28.70	61.16	21.88	28.71
18	41.40	31.53	28.56	41.40	15	25.73	38.92	21.40	25.73
16	41.31	57.13	28.03	41.31	19	24.69	4.39	21.31	24.70
22	41.05	85.83	28.21	41.06	44	23.42	7.15	21.08	23.42
47	40.78	82.33	27.91	40.78	9	22.88	96.22	21.22	22.89
61	40.22	12.20	28.44	40.22	33	22.36	29.89	19.50	22.36
40	40.01	84.98	28.38	40.02	39	20.69	57.73	19.26	20.69
50	39.35	56.04	28.64	39.35	69	20.39	63.02	19.26	20.40
27	38.77	34.07	27.87	38.77	3	17.12	28.70	18.66	17.13
48	38.35	23.77	27.57	38.35	24	16.96	42.73	18.68	16.98
34	38.29	5.69	28.08	38.30	68	14.53	13.63	7.69	14.53
12	37.96	85.14	27.44	37.96	60	13.66	22.69	8.15	13.66

Table 7: Comparisons between different data augmentation methods on the span-level exact F₁ score of every crowd worker. **Ori.** stands for the original score in real datasets before any augmentation. **Rnd. Gen.** is a naive augmentation method with random generated annotations. **SES Only** indicates the *shifting*, *shrinking*, and *expanding* method we proposed. **SES + Alg.2** means SES with Algorithm 2 which is our final method.