
Student as an Inherent Denoiser of Noisy Teacher

Jiachen Zhao

University of Massachusetts Amherst
jiachenzhao@umass.edu

Abstract

Knowledge distillation (KD) has been widely employed to transfer knowledge from a large language model (LLM) to a specialized model in low-data regimes through pseudo label learning. However, pseudo labels generated by teacher models are usually noisy and may influence KD performance. This study delves into KD with noisy teachers and uncovers that the student model can already generate more accurate predictions than the teacher labels used to train it in the middle of KD, indicating its inherent ability to *denoise* noisy teacher labels. Motivated by this finding, we propose Peer-Advised KD to improve vanilla KD from noisy teachers. Experiments show that Peer-Advised KD can outperform LLM by approximately 5% with 50 human-labeled data, and even competitive to standard supervised finetuning with 750 human-labeled data.

1 Introduction

Recently, with the emergence of large language models (LLMs), knowledge distillation (KD) with label matching has been widely applied to low-data regimes (28; 12; 17; 10; 33; 37; 7). In scenarios where there is a scarcity of labeled in-domain data but an abundance of unlabeled data, LLMs can serve as an annotator to generate pseudo labels through few-shot prompting, i.e., in-context learning (2). These pseudo labels are then used to train a smaller student model through KD. The student model can offer improved efficiency and faster inference for deployment. However, one crucial problem with such use of KD is that the generated pseudo labels of teacher models are usually noisy. The performance of student models can be greatly dependent on the quality of teacher labels (13). Despite many past works on leveraging LLMs as annotators, few works have investigated how to optimize learning noisy data. Some previous works for noisy label learning utilize loss reweighting (11; 13; 16) that may not apply to few-shot cases where pseudo labels are overwhelmingly more than gold labels, leading to great imbalance and unstable training.

Therefore, this work aims at providing new insights into learning noisy teacher labels for more reliable student training in low-data regimes. This work mainly focuses on learning linguistic structures as our test bed, which usually requires extensive labeled data to acquire (8) and is the fundamental ability to downstream NLP tasks (20). More specifically, we first analyze the learning process of KD (Sec. 2.1) with according empirical results (Sec. 2.2). Our analysis reveals that **during KD, the student model exhibits a denoising ability, generating superior predictions compared to the noisy teacher labels used to train it.** Our findings imply that vanilla KD from noisy teacher gives sub-optimal student’s performance. Motivated by our findings, we present Peer-Advised KD (PA-KD) in Sec. 3 which harnesses the student’s inherent denoising ability to generate improved pseudo labels for KD.

2 Revisiting KD

Definition 1 (KD for Semi-Supervised Learning) *We assume there is an accessible set of labeled data $\mathcal{D}_{(x,y)}^{\text{labeled}}$ whose distribution is p^* . Then, given a set of unlabeled data $\mathcal{D}_x^{\text{unlabeled}}$ sampled from the same distribution p^* , T is employed to generate the distillation set $\mathcal{D}_{(x,\tilde{y})}^{\text{pseudo}}$. The student model S is then trained by minimizing the loss $\mathcal{L} = \beta\ell(\hat{y}, y) + \ell(\hat{y}, \tilde{y})$.*

We first detail the setting of KD in our analysis. The student model is denoted by S whose label is \hat{y} , and the teacher model is denoted by T whose label is \tilde{y} . The teacher model T may leverage $\mathcal{D}_{(x,y)}^{\text{labeled}}$ through In-Context Learning (2) or supervised finetuning to generate pseudo labels. Additionally, we set $\beta=0$ to focus on the effects of distillation set following (29).

2.1 Decoding the Learning Process in KD

This section analyzes the learning process during KD theoretically. We hypothesize that at the beginning of KD, the student S will first fully converge to partial pseudo labels of teacher T (the set of which is denoted as Y_{learnt} , while the corresponding x is in X_{learnt}). Because the learning process is usually imbalanced where the weights are updated more frequently for information that occurs more often in the training dataset (30; 19; 5; 35). Formally, after t_0 training steps, we consider S has learnt a set of knowledge denoted by ξ_{base} and has nearly zero loss on Y_{learnt} . At the next training step, we can denote $\hat{y} = S(x|\xi_{\text{base}})$ which means the student’s prediction is conditioned on its learnt knowledge ξ_{base} . Accordingly, we represent teacher’s encoded knowledge with ξ_T . Then, for the next incoming training datapoint (x, \tilde{y}) , we consider two cases depending on underlying patterns of x and analyze respective learning effects. We focus on the comparison between student’s predictions and teacher labels to learn. Formally, we define $\delta := s(y, S(x|\xi_{\text{base}})) - s(y, \tilde{y})$ where s is a score function (F1 score for parsing) for similarity measurement.

Case 1: the underlying patterns of x are the same as or similar to patterns of instances in X_{learnt} . (1) If $\tilde{y} \in Y_{\text{learnt}}$, i.e., teacher labels that S has converged to, then we can get $\delta = 0$. (2) If $\tilde{y} \notin Y_{\text{learnt}}$, we speculate that δ can be unclear. δ may be positive (i.e., $s(y, S(x|\xi_{\text{base}})) > s(y, \tilde{y})$), indicating the student generates better labels than the teacher. This is because ξ_{base} might be much richer than ξ_T for the specific pattern of x , especially when $|\mathcal{D}_x^{\text{unlabeled}}| \gg |\mathcal{D}_{(x,y)}^{\text{labeled}}|$. Chances are that $(X_{\text{learnt}}, Y_{\text{learnt}})$ after some t_0 may already contain extensive credible information about diverse cases similar to x , contributing to strong ξ_{base} which is conditioned on by S for its prediction. In this scenario during KD, the student’s performance may be degraded because it is forced to learn inferior teacher labels to its own predictions. We empirically demonstrate this scenario in Sec. 2.2, which we refer to as student’s inherent denoising ability.

Case 2: the underlying patterns of x are clearly distinct from patterns of instances in X_{learnt} . (1) δ may tend to be negative indicating that the teacher labels are more correct (i.e., $s(y, S(x|\xi_{\text{base}})) < s(y, \tilde{y})$). Because knowledge of ξ_{base} involving patterns of x may have not been learnt effectively by S after t_0 . In this case, S may benefit from KD by learning new knowledge. (2) If patterns of x are uncommon, δ can be unclear because both student and teacher may be making educated guesses conditioned on learnt knowledge. Chances are that both predictions are not credible.

All in all, we hypothesize that KD is an iterative process of alternating between case 1 and case 2 discussed above. This analysis of KD can be viewed as “expansion” (36) from the learnt information to the under-learnt.

2.2 Student Can Inherently Denoise Teacher Labels

In this section, experiments are conducted to provide empirical understandings of KD. We employ two datasets Penn Tree Bank and CRAFT. Teacher models are respectively Codex (3) and T5-base (24). F1 score is used to compare the similarity of labels (i.e., parsed trees). Detailed implementations are in Appendix D.2. During knowledge distillation (KD), the student’s predictions on training data are compared with the teacher labels used for training it. Notably, we observe that the student demonstrates a denoising ability, generating better labels than some teachers labels used to train it in the middle of KD as shown in Fig. 1. This implies vanilla KD is sub-optimal since the student is forced to learn inferior teacher labels to its own predictions sometimes.




Figure 1: During KD, the student can actually generate better labels than the teacher labels used to train it. The y-axis is percentage that the student’s predicted labels are better than teacher’s pseudo labels for training (see definition for δ in Sec. 2.1). The x-axis is the same as Fig. 2.




Figure 2: The student model tends to converge more to less noisy pseudo labels of teacher during training. The y-axis is F1 w.r.t ground truth. The x-axis is sets of data ranked in descending order based on student’s convergence with teacher labels. Each interval range is 5% of the size of dataset. From Fig. 2a to Fig. 2c, the general noise level of teacher labels is decreasing.

How does the denoising effect occur? We identify such denoising effect can be attributed to the discrepancies in learning teacher labels, as shown in Fig. 2. Student converges faster to less noisy teacher labels than noisier ones. Arpit et al. (1) observed such discrepancies in adversarial label learning where labels are randomly flipped to introduce noises. Our results suggest that the finding also holds for sequence-level KD where noises in labels are not random but generated by a teacher model, implying its inductive bias. Further analysis is provided in Appendix A. When examining Fig. 2a and Fig. 1a (or Fig. 2b and Fig. 1b, or Fig. 2c and Fig. 1c), the student is observed to, after converging to more credible teacher labels (higher ground-truth F1), outperform the noisier teacher labels (lower ground-truth F1) by conditioning on its learnt knowledge. This falls into the scenario (2) in the Case 1 of our theoretical framework of KD in Sec. 2.1. Due to student’s inclination to converging to cleaner teacher labels, the student’s learnt knowledge ξ_{base} at some time in the middle of KD can be possibly reliable and even much more sufficient than teacher’s ξ_T for some patterns, especially when $|\mathcal{D}_{(x,x)}^{unlabeled}| \gg |\mathcal{D}_{(x,y)}^{labeled}|$. This is also evidenced by experiments in Appendix B showing that denoising ability is more noticeable when the distillation set (i.e., noisy teacher’s labels) is much larger than the real labels used to train the teacher model. Further analysis between our theoretical framework of KD and empirical results are shown in Appendix C.

3 Peer-Advised Knowledge Distillation from Noisy Teachers

In this section, we propose Peer-Advised Knowledge Distillation (PA-KD) for learning from noisy teachers in low-data regimes motivated by our findings. We have shown that during KD, the student can outperform the noisier teacher labels by conditioning on its learnt knowledge because the student tends to converge to more credible teacher labels faster. To leverage and magnify this process in KD, we first extract teacher labels denoted by \mathcal{T}_{high} based on convergence to train a peer student model S_1 , which is then employed to re-annotate the left teacher labels other than \mathcal{T}_{high} to get peer labels \mathcal{S} . More specifically, within dataset $\mathcal{T} := \{(x, \tilde{y})\}$ and the set of all teacher labels \tilde{Y} , a partition

	#Params	PTB		CRAFT	
# Human-labeled Data		50	250	50	250
Teacher: Code-Davinci-002	175B	74.51	76.49	62.27	64.37
Student: T5-base	220M	0	58.02	0	54.11
SLKD	220M	77.21	77.40	65.58	66.71
SLKD+SD	220M	77.34	77.78	67.21	68.05
SLKD+SD w/ HC	220M	77.11	78.22	67.33	68.15
Selective KD	220M	78.13	78.44	68.24	68.37
PA-KD	220M	78.64	78.98	68.98	68.92

Table 1: Performance (F1 score (%)) of different methods on PTB and CRAFT in two few-shot settings.

function is defined as follow,

$$\mathcal{T}_{\text{high}} = \{(x, \tilde{y}) | s(\tilde{y}, \hat{y}) \geq M(\tilde{Y}, r)\}, \quad (1)$$

where $M(\tilde{Y}, r)$ stands for the threshold score of top $r\%$ predictions among \tilde{Y} and s is F1 score function to measure the similarity between the student labels \hat{y} and teacher labels \tilde{y} as a proxy indicating convergence. Finally, the ultimate student model S_2 will be trained with $\mathcal{T}_{\text{high}}$ and \mathcal{S} for improved performance. In addition, we propose another baseline, i.e., selective KD, which uses $\mathcal{T}_{\text{high}}$ alone for KD rather than using full data to reduce influence of noises in teacher labels.

3.1 Experimental Results

We utilize the Penn Treebank (PTB) corpus (21) and the Colorado Richly Annotated Full-Text (CRAFT) corpus (4) in our experiments for grammar induction. More details on the dataset and implementations are shown in Appendix D.1 and Appendix D.3 respectively. We consider three baselines for comparison with our proposed PA-KD. (1) SLKD: knowledge distillation by directly learning decoded sequences proposed by Kim and Rush (14), which is the default way to do KD throughout this paper. (2) SLKD+SD: after KD, the student is retrained through self-distillation (SD) (9; 23; 18), where labels of the same training data generated by itself are employed. (3) SLKD+SD w/HC: A variant of baseline 2 is to conduct SD only with labels that have high output confidence (HC) of the model, which is widely adopted to obtain better pseudo labels (31; 32; 22; 15).

Results are displayed in Table 1. With few-shot prompting, the LLM with limited labeled data can give decent performance. In comparison, the T5-base model trained with a small set of labeled data performs much worse than teacher. When there are 50 real data, T5 base cannot produce meaningful predictions. However, when training T5 base with massive pseudo labels through KD, it can outperform the teacher. Furthermore, our Peer-Advised KD can significantly improve vanilla KD and surpass all baselines. For example, on PTB, PA-KD improves student’s F1 by around 2%. Additionally, PA-KD with 50 labeled data can outperform finetuning the student model with 500 data (results shown in Fig. 3 in Appendix). PA-KD can thus greatly save labeling cost. On the other hand, Table 1 demonstrates that selective KD using highly converged teacher labels achieves impressive performance, surpassing KD on the entire dataset. This suggests that vanilla KD on full data is sub-optimal due to the noise present in the teacher labels. By utilizing our finding that student tends to converge faster to cleaner teacher labels, noisier teacher labels may be filtered based on convergence and thus student can avoid learning misleading noise, resulting in improved KD performance but with much fewer data. Our approach is also effective for supervised teacher and self-distillation (see Appendix E and Appendix F).

4 Conclusion

In this work, we investigate knowledge distillation from a noisy teacher (e.g., LLM) through pseudo label learning. Extensive experiments show that the student inclines to converge to cleaner teacher labels and can utilize its acquired knowledge to denoise noisy teacher labels during KD. Motivated by our findings, we then propose a novel Peer-Advised KD approach which leverages highly-converged

teacher labels to train a peer model to relabel under-converged teacher labels (that are possibly noisier). Experiments show our method is label-efficient and can consistently outperform strong baselines on two public datasets in low-data regimes.

References

- [1] Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International conference on machine learning*, pages 233–242. PMLR, 2017.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [3] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- [4] K Bretonnel Cohen, Karin Verspoor, Karën Fort, Christopher Funk, Michael Bada, Martha Palmer, and Lawrence E Hunter. The colorado richly annotated full text (craft) corpus: Multi-model annotation in the biomedical domain. In *Handbook of Linguistic Annotation*, pages 1379–1394. Springer, 2017.
- [5] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019.
- [6] Mostafa Dehghani, Arash Mehrjou, Stephan Gouws, Jaap Kamps, and Bernhard Schölkopf. Fidelity-weighted learning. *arXiv preprint arXiv:1711.02799*, 2017.
- [7] Bosheng Ding, Chengwei Qin, Linlin Liu, Lidong Bing, Shafiq Joty, and Boyang Li. Is gpt-3 a good data annotator? *arXiv preprint arXiv:2212.10450*, 2022.
- [8] Andrew Drozdov, Patrick Verga, Mohit Yadav, Mohit Iyyer, and Andrew McCallum. Unsupervised latent tree induction with deep inside-outside recursive auto-encoders. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1129–1141. Association for Computational Linguistics, June 2019.
- [9] Tommaso Furlanello, Zachary Chase Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *International Conference on Machine Learning*, 2018.
- [10] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. Chatgpt outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv:2303.15056*, 2023.
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [12] Namgyu Ho, Laura Schmid, and Se-Young Yun. Large language models are reasoning teachers. *arXiv preprint arXiv:2212.10071*, 2022.
- [13] Fotis Iliopoulos, Vasilis Kontonis, Cenk Baykal, Gaurav Menghani, Khoa Trinh, and Erik Vee. Weighted distillation with unlabeled examples. In *Advances in Neural Information Processing Systems*, volume 35, 2022.
- [14] Yoon Kim and Alexander M. Rush. Sequence-level knowledge distillation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327. Association for Computational Linguistics, November 2016.

- [15] Hunter Lang, Monica N Agrawal, Yoon Kim, and David Sontag. Co-training improves prompt-based learning for large language models. In *International Conference on Machine Learning*, pages 11985–12003. PMLR, 2022.
- [16] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013.
- [17] Zongjie Li, Chaozheng Wang, Pingchuan Ma, Chaowei Liu, Shuai Wang, Daoyuan Wu, and Cuiyun Gao. On the feasibility of specialized ability stealing for large language code models. *arXiv preprint arXiv:2303.03012*, 2023.
- [18] Yang Liu, Sheng Shen, and Mirella Lapata. Noisy self-knowledge distillation for text summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 692–703. Association for Computational Linguistics, June 2021.
- [19] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2537–2546, 2019.
- [20] Christopher D Manning, Kevin Clark, John Hewitt, Urvashi Khandelwal, and Omer Levy. Emergent linguistic structure in artificial neural networks trained by self-supervision. *Proceedings of the National Academy of Sciences*, 117(48):30046–30054, 2020.
- [21] Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330, 1993.
- [22] David McClosky, Eugene Charniak, and Mark Johnson. Effective self-training for parsing. In *Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 4-9, 2006*. The Association for Computational Linguistics, 2006.
- [23] Hossein Mobahi, Mehrdad Farajtabar, and Peter L. Bartlett. Self-distillation amplifies regularization in hilbert space. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20, Red Hook, NY, USA, 2020*. Curran Associates Inc. ISBN 9781713829546.
- [24] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [25] Haoyue Shi, Karen Livescu, and Kevin Gimpel. On the role of supervision in unsupervised constituency parsing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, November 16-20, 2020*, pages 7611–7621. Association for Computational Linguistics, 2020.
- [26] Richard Shin and Benjamin Van Durme. Few-shot semantic parsing with language models trained on code. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, July 10-15, 2022*, pages 5417–5425. Association for Computational Linguistics, 2022.
- [27] Richard Shin, Christopher H. Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. Constrained language models yield few-shot semantic parsers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, 7-11 November, 2021*, pages 7699–7715. Association for Computational Linguistics, 2021.
- [28] Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. Distilling multi-step reasoning capabilities of large language models into smaller models via semantic decompositions. *arXiv preprint arXiv:2212.00193*, 2022.

- [29] Samuel Stanton, Pavel Izmailov, Polina Kirichenko, Alexander A Alemi, and Andrew G Wilson. Does knowledge distillation really work? *Advances in Neural Information Processing Systems*, 34:6906–6919, 2021.
- [30] Yanmin Sun, Andrew KC Wong, and Mohamed S Kamel. Classification of imbalanced data: A review. *International journal of pattern recognition and artificial intelligence*, 23(04):687–719, 2009.
- [31] Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. Grammar as a foreign language. *Advances in neural information processing systems*, 28, 2015.
- [32] Fusheng Wang, Jianhao Yan, Fandong Meng, and Jie Zhou. Selective knowledge distillation for neural machine translation. 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 6456–6466, 2021.
- [33] Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. Want to reduce labeling cost? GPT-3 can help. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics, November 2021.
- [34] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 322–330, 2019.
- [35] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. *Advances in neural information processing systems*, 30, 2017.
- [36] Colin Wei, Kendrick Shen, Yining Chen, and Tengyu Ma. Theoretical analysis of self-training with deep networks on unlabeled data. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=rC8sJ4i6kaH>.
- [37] Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. GPT3Mix: Leveraging large-scale language models for text augmentation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. Association for Computational Linguistics, November 2021.




Figure 3: Comparison between Peer-Advised KD (PA-KD) with 50 labeled data and supervised-finetuning (SFT) T5-base with different amounts of labeled data.




Figure 4: Convergence comparison (indicated by F1 w.r.t teacher labels) for students trained with good teacher labels (S_{High}) and bad teacher labels (S_{Low}). The division of teacher labels is binary and based on ground-truth F1 scores, e.g., the half with higher F1 w.r.t ground-truth is used to train S_{High} .

A Discrepancy in Convergence

We find during KD, the student demonstrates a greater inclination towards aligning with high-quality teacher labels. In all cases of Fig. 2, as the quality of teacher labels (measured by ground-truth F1) decreases, we can observe a gradual reduction in the student’s convergence with the teacher labels. Fig. 4 clearly shows the student tends to converge faster to good teacher labels than bad ones. Such disparity in convergence has also been observed in other domains (6; 34). Thus, noise level of teacher labels can be estimated by convergence extent of student with teacher, which may be leveraged for label filtering. In terms of why less noisy teacher pseudo labels are more learnable, we speculate that the main reason is the inconsistency of errors in teacher labels. The errors made by teacher can be arbitrary and instance-specific. Such lack of consistent patterns across noisy labels prevents student from applying what it has learnt from one training sample to another, thus leading to slow convergence in learning.

B Denoising Effects and the Size of Distillation Sets

We identify that the emergence of denoising effect is related to the size of distillation set. As shown in Fig. 5, the average percentage that the student denoise teacher’s labels will increase with more teacher labels used for training. When the number of teacher labels for student training is much larger than data (500 real labels) used to train teacher model, the student can have noticeable denoising effect despite the smaller size. This may support our hypothesized situation (2) in Case 1 of Sec. 2.1 that sufficient teacher labels can contain more information than a small set of real labels.




Figure 5: The figure shows the change of average percentage that student outperforms its teacher during KD with respect to the size of distillation set on PTB. In this case, the student is T5-small, while the teacher is T5-base trained with 500 real labels.

C Further Analysis of Denoising Effects of Student

Denoising effects may diminish as convergence decreases, particularly at the early epoch (see Fig. 1b and Fig. 1a) and for the case that the teacher is trained through supervised tuning in Fig. 1c). We attribute this observation to the first scenario in Case 2 in Sec. 2.1. The reduced denoise effect of student can be caused by the increase of unfamiliar knowledge involved in samples that have low convergence between student and teacher labels. Additionally, random spikes are observed in cases of extremely low convergence (see Fig. 1b and Fig. 1c). In those cases, both the student and teacher actually exhibit low ground-truth F1 scores (see Fig. 2b and Fig. 2c), indicating that they might be making educated guesses given some rare patterns or features in the input data. This may be an example of our second hypothesized scenario in Case 2 in Sec. 2.1

D Experiments Details

D.1 Datasets

We process PTB dataset with the standard splits (2-21 for training containing 39832 sentences, 22 for validation, 23 for test). We randomly split the CRAFT corpus into a training set, a validation set and a test set following the ratio of 6:2:2. We consider two low-data regimes. We randomly select 50 and 250 samples in the training set as two cases of available annotated data, while the rest are treated as unlabeled data.

D.2 Implementations for Experiments in Sec. 2.2

We assume only 250 real labeled data available, while the rest of training data are considered as unlabeled. We employ Codex-Davinci-002 as teacher model and T5-base as student. We use 50 exemplars for PTB and 35 exemplars for CRAFT in terms of prompting LLM. On the other hand, for KD with a supervised teacher, the teacher model is T5-base and the student is T5-small. The teacher is fine-tuned with 500 in-domain data for 15 epochs on PTB. All supervised training entails a learning rate of 3×10^{-4} and a batch size of 32.

D.3 Implementations for Experiments in Sec. 3.1

In knowledge distillation, the student model is T5 (24) trained using hard labels, i.e., decoded sequences, also known as sequence-level KD (14). For prompting, we utilize Codex (code-davinci-002) as the teacher model which is shown effective for parsing (27; 26). To avoid the number of demonstrations from surpassing token limits, we use 50 exemplars for PTB and 35 exemplars for CRAFT. Supervised training entails a learning rate of 3×10^{-4} , a batch size of 32, and 20 epochs. For convergence-based partition in Peer-Advised KD, an initial student S_0 is trained with 2 epochs whose predictions on training data are compared with teacher labels for deciding convergence. r is set as 50%, i.e., the half of teacher labels with greater convergence with student’s predictions are considered as $\mathcal{T}_{\text{high}}$. For SD w/ HC, weights with highest validation score is used to generated labels, while for

sole SD, student is trained for 4 epochs to avoid generating the same labels to teacher labels due to overfitting. We set the half with above-mean output probabilities as HC labels. During evaluation, sentence-level unlabeled parsing F1 scores are reported. These scores are computed individually for each sentence and then averaged across the dataset, aligning with previous works on parsing (25).

E Results on Peer-Advised KD with Supervised Teacher

Our proposed peer-advised KD approach also works for the supervised teacher. We use T5-base as teacher model and T5-small as the student and assume access to 500 human-labeled data. Results are shown in Table 2. In this case, a much smaller student (T5-small) trained with real data can only reach 21.92% F1 on PTB, which is not comparable to the teacher, T5-base. KD can still significantly improve T5-small’s performance, but the student still performs inferior to the teacher due to incapacity. However, our peer-advised KD surpasses the vanilla KD by 1.94%, enabling the student to be competitive to the teacher.

Methods	#Params	F1
Teacher: T5-base	220M	74.21
Student: T5-small	60M	21.92
SLKD	60M	73.05
SLKD+SD	60M	74.12
SLKD+SD w/HC	60M	71.17
Selective KD	60M	73.72
PA-KD	60M	74.60

Table 2: Experimental results on knowledge distillation from a supervised-trained teacher.

F Results on Peer-Advised KD for Self-distillation

Methods	F1
T5-base	74.21
SD	79.17
SD w/HC	78.89
Selective SD	81.02
PA-KD	82.56

Table 3: Experimental results of self-distillation (SD) on PTB.

In addition to employing a capable teacher model, we explore a common scenario where only a single model is accessible. In this scenario, self-distillation (SD) (9; 23; 18) proves to be valuable for enhancing generalization by training the model with self-generated labels. Notably, our proposed Peer-Advised KD is applicable to self-distillation as well. The results in Table 3 demonstrate that the single-peer PA-KD significantly enhances the performance of the supervised model by approximately 8% F1, outperforming both direct SD and SD with high-confidence labels.