

PseudoReasoner: Leveraging Pseudo Labels for Commonsense Knowledge Base Population

Tianqing Fang¹, Quyet V. Do¹, Hongming Zhang², Yangqiu Song¹
Ginny Y. Wong³, Simon See³,

¹Department of Computer Science and Engineering, HKUST, Hong Kong SAR

²Tencent AI Lab, Bellevue, USA

³NVIDIA AI Technology Center (NVAITC), NVIDIA, Santa Clara, USA

{tfangaa, vqdo, yqsong}@cse.ust.hk

hongmzhang@global.tencent.com, {gwong, ssee}@nvidia.com

Abstract

Commonsense Knowledge Base (CSKB) Population aims at reasoning over unseen entities and assertions on CSKBs, and is an important yet hard commonsense reasoning task. One challenge is that it requires out-of-domain generalization ability as the source CSKB for training is of a relatively smaller scale (1M) while the whole candidate space for population is way larger (200M). We propose PseudoReasoner, a semi-supervised learning framework for CSKB population that uses a teacher model pre-trained on CSKBs to provide pseudo labels on the unlabeled candidate dataset for a student model to learn from. The teacher can be a generative model rather than restricted to discriminative models as previous works. In addition, we design a new filtering procedure for pseudo labels based on influence function and the student model's prediction to further improve the performance. The framework can improve the backbone model KG-BERT (RoBERTa-large) by 3.3 points on the overall performance and especially, 5.3 points on the out-of-domain performance, and achieves the state-of-the-art. Codes and data are available at <https://github.com/HKUST-KnowComp/PseudoReasoner>.

1 Introduction

Commonsense knowledge are the common agreements by most people on daily entities, which are crucial for intelligent systems to act sensibly in the real world (Davis, 1990; Liu and Singh, 2004). Endowing natural language understanding systems with the ability to draw commonsense reasoning remains an important yet challenging task.

Throughout the development of automated commonsense understanding, CommonSense Knowledge Base (CSKB) is an important form of automatic commonsense reasoning system to store knowledge sources for drawing inferences. With expert-curated relations and human annotations, CSKBs such as ConceptNet (Liu and Singh, 2004),

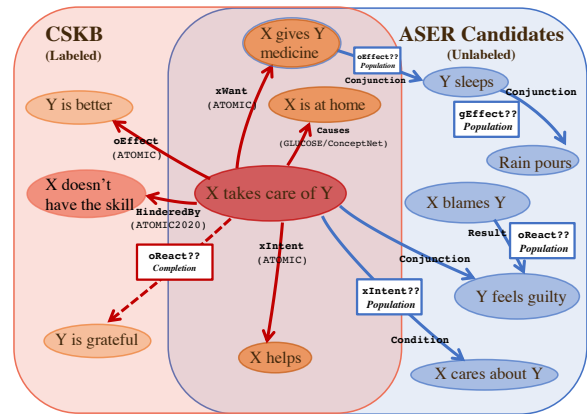


Figure 1: An example of CSKB Population. The coral part (left) and the blue part (right) respectively represent the labeled CSKBs and the unlabeled candidate pool. The entities in the overlap parts are marked with coral shape and blue outline. The reasoning within the CSKB (coral-outlined boxes) belongs to the CSKB Completion part, and the reasoning that are not limited within the domain of CSKBs belongs to CSKB Population.

ATOMIC (Sap et al., 2019; Hwang et al., 2021), and GLUCOSE (Mostafazadeh et al., 2020) are developed to study commonsense regarding properties of objects, causes and effects of events and activities, motivations and emotional trajectories of humans on certain circumstances, and so on. As those human-annotated CSKBs are sparse and usually of a small scale and coverage, reasoning tasks on CSKB such as *CSKB Completion* (Li et al., 2016; Saito et al., 2018; Malaviya et al., 2020) and *CSKB Population* (Fang et al., 2021b,a) are defined with the goal of either adding new edges/assertions within the training knowledge base (*CSKB Completion*), or adding new edges/assertions from outside of CSKBs (*CSKB Population*). A visualized comparison between the two tasks is shown in Figure 1.

Different from CSKB Completion, which adopts a *close-world assumption* and assume all knowledge are in-domain, the population task deals with unseen entities and requires a more out-of-distribution reasoning ability. In this paper, we

study commonsense reasoning in the context of CSKB Population. In this task, four mainstream CSKBs, ConceptNet, ATOMIC, ATOMIC₂₀²⁰, and GLUCOSE are aligned together as the labeled dataset. ASER (Zhang et al., 2022), a large-scale eventuality (events, activities, and states) knowledge graph is aligned with CSKBs and serves as unlabeled candidates for populating commonsense knowledge. Human annotations on held-out dev/test sets sampled from both CSKBs and ASER are provided as the evaluation set.

There are two major challenges remaining unsolved for CSKB Population. First, the scale of the annotated training set (ConceptNet, ATOMIC, and GLUCOSE) is approximately 1M samples, too small compared with 200M of the actual candidate space to perform population (ASER). Second, as inherently only ground-truth (positive) examples are provided by CSKBs, the randomly sampled negative examples in the task are less informative and may lead the model to overfit artifacts of the dataset. A supervised learning model finetuned on such annotated training set is hard to be generalized to out-of-domain knowledge space, as shown in Fang et al. (2021a) and also Table 3 in our paper, where the AUC for out-of-domain test sets performs over 10 points worse than the in-domain part.

To address the above challenges, we propose **PseudoReasoner**, a semi-supervised learning framework that uses a pre-trained commonsense teacher model to automatically label the unlabeled candidates to serve as pseudo labels, such that the student model can be further finetuned with pseudo labels to improve out-of-domain commonsense reasoning ability. In fact, pre-trained language models finetuned on commonsense knowledge bases have shown to perform generalizable commonsense reasoning on downstream tasks to some extent. For example, leveraging commonsense knowledge generated by COMET (Bosselut et al., 2019), a language model finetuned on ATOMIC, can improve the performance on commonsense QA (Yang et al., 2020; Schwartz et al., 2020; Bosselut et al., 2021). Different from the text generation paradigm as in previous works, here we leverage the commonsense language model as a teacher model for labeling unlabeled candidates. To further improve the quality of pseudo labels, we use both influence function (Koh and Liang, 2017) and the student model’s prediction to select highly confident pseudo exam-

ples.

Our contribution is three-fold:

- 1) We introduce a new way of providing pseudo label for CSKB Population by leveraging generative commonsense language models.
- 2) We propose a semi-supervised learning framework with pseudo labels and a special filtering mechanism based on influence function and student model’s prediction that significantly improve the performance of CSKB Population, especially for out-of-domain knowledge triples.
- 3) We demonstrate the effectiveness of our framework by extensive experiments on different backbone models and different semi-supervised learning methods. We achieve the state-of-the-art performance on this task.

2 Related Works

2.1 Commonsense Reasoning over Knowledge Bases

Commonsense Knowledge Bases (CSKBs) provides rich human-curated commonsense knowledge in the form of head-relation-tail triples and are perfect fields to conduct commonsense reasoning. Reasoning tasks defined over CSKBs can be categorized into two main categories, CSKB Completion and CSKB Population. CSKB Completion adopts a close-world assumption and assume the knowledge not included in the knowledge base are incorrect, thus using evaluation metrics such as accuracy, MRR, and HITS@k (under the link prediction setting (Li et al., 2016; Malaviya et al., 2020; Saito et al., 2018; Jastrzēbski et al., 2018)) or BLEU (under the text generation setting (Bosselut et al., 2019; Hwang et al., 2021)) to evaluate the reasoning performance. As vanilla knowledge base completion models such as TransE (Bordes et al., 2013) and ComplEx (Trouillon et al., 2016) do not take the node semantics into account and CSKBs are much sparser than standard factual KBs, Malaviya et al. (2020) and Wang et al. (2021) use BERT (Devlin et al., 2019) to encode nodes and propose graph densifiers to address the sparsity issue.

CSKB Population, on the other hand, requires the model to reason on not only existing nodes in the CSKB (Fang et al., 2021b,a), but the nodes from other domains (e.g., ASER (Zhang et al., 2022)). The evaluation set is manually annotated to form a binary classification task, and AUC is used as

	D_l	D_u
# triples	1,119,517 (training)	218,809,746

Table 1: Statistics of labeled D_l (CSKB with negative examples) and unlabeled D_u (processed ASER).

the evaluation metric. The population task is inherently a classification task, KG-BERT (Yao et al., 2019) and KG-BERTSage (Fang et al., 2021a), a graph-aware model based on KG-BERT, are used to tackle the task. While there are other ways of populating commonsense knowledge, such as leveraging conceptualization (He et al., 2022) or generative language models (Bosselut et al., 2021; West et al., 2022), they are different with CSKB Population in several aspects. For conceptualization, they provide new triples in depth through hierarchical structures of concepts by lexical manipulation and the core semantics are preserved, while we do population in breadth by exploring more diverse information-extracted triples. For text generation, first, text generation itself doesn’t perform reasoning, while discrimination on a knowledge graph can provide structural and contextual information to do reasoning. Second, generative models are trained only on positive/plausible examples by nature while do not handle negative/implausible examples. We will show the importance of negative examples in CSKB Population in the experiments.

2.2 Pseudo Labels

Pseudo labels are widely used in the semi-supervised learning setting (Isken et al., 2019; Xie et al., 2020b; Sohn et al., 2020; Pham et al., 2021). In common, pseudo labels alleviates the cost of human annotation, or familiarizes the student model with out-of-domain or same-distribution-but-unseen data (i.e data from the large unlabeled dataset) under the advice from teacher models. It has been used in image classification (Lee et al., 2013; Shi et al., 2018; Yalniz et al., 2019; Xie et al., 2020b; Pham et al., 2021), machine translation (He et al., 2020; Chen et al., 2021), information retrieval (Wang et al., 2022), and word segmentation (Huang et al., 2021). In the simplest form of pseudo labels, the teacher model that provides pseudo labels is static or iteratively updated using the latest student model. Meta pseudo label (Pham et al., 2021), on the other hand, leverages meta-learning technique to learn an end-to-end teacher-student network that updates both networks jointly based on the performance of the student network on the labeled dataset, yielding state-of-the-art per-

	<i>Original Test Set</i>	<i>CSKB head + ASER tail</i>	<i>ASER edges</i>
Sampled from	D_l	D_l head, D_u tail	D_u
Domain	In-domain	Out-of-domain	Out-of-domain
# triples (dev)	2,042	2,193	1,982
# triples (test)	8,437	9,103	7,974

Table 2: The details of evaluation set categorization.

formance on ImageNet (Deng et al., 2009).

2.3 Other Semi-supervised Learning Methods

Most works on semi-supervised learning leverages the idea of consistency training (Zhou et al., 2003; Xie et al., 2020a; Gururangan et al., 2019), which aims to constrain the model to be robust given noised inputs or hidden states. UDA (Unsupervised Data Augmentation) (Xie et al., 2020a) adopts back translation and TF-IDF word replacement to provide noisy inputs and use them for consistency training. However, in the setting of CSKB Population, the problem of out-of-distribution is more important as we aim to explore novel commonsense knowledge that are not seen in the training domain, which makes directly finetuning on pseudo labels more effective than using consistency training.

3 CSKB Population

3.1 Problem Definition

Denote a labeled ground-truth Commonsense Knowledge Base as $D_l^+ = \{(h, r, t) | h \in H_C, r \in R, t \in T_C\}$. The overall labeled dataset $D_l = \{(h, r, t), y\}$ is composed of D_l^+ and a randomly sampled negative dataset D_l^- from the CSKB, where $y \in \{0, 1\}$ is the label of the triple. Triples from D_l^+ are labeled 1 while those from D_l^- are labeled 0. H_C and T_C are the set of heads and tails in the CSKB. R is the relation set.

$D_u = \{(h, r, t) | h \in H_u, r \in R, t \in T_u\}$ is an unlabeled candidate knowledge base of the same format and relation set as the CSKB. It is of a much larger scale than the labeled part and is a source for populating commonsense knowledge. H_u and T_u is the set of heads and tails in the unlabeled KB. The task is defined as given the labeled commonsense knowledge base D_l as training source, predict the plausibility of triples from D_u .

We use the dataset provided by Fang et al. (2021a). Here, the set of heads H_C , tails T_C , and relations R come from the alignment of ConceptNet, ATOMIC, ATOMIC₂₀ (newly developed relations), and GLUCOSE, as in the original paper. The unlabeled KB D_u is adapted from ASER, where the discourse relations are converted to commonsense

relations to serve as candidates for population. The evaluation dataset with 32K triples is sampled from both D_l and D_u and manually annotated. There are three categories of the evaluation set, *Original Test Set*, *CSKB head + ASER tail*, and *ASER edges*, where the first category is sampled from the held-out test split in D_l (both D_l^+ and negative examples D_l^-) and is thus an in-domain test set, and the latter two are novel assertions outside of D_l and are thus out-of-domain. The statistics and descriptions of the training and evaluation datasets are shown in Table 1 and 2. The list of commonsense relations, alignment methods, and more detailed statistics of the dataset can be found in the Appendix A.

3.2 Backbone Models

Considering the nature of the CSKB Population task is triple classification in the form of natural language, we use KG-BERT (Yao et al., 2019) as the backbone model following Fang et al. (2021a). In detail, a triple (h, r, t) is concatenated and serialized as “[CLS], $h_1, \dots, h_{|h|}$, [SEP], $[r]$, [SEP], $t_1, \dots, t_{|t|}$ ”. Here, [CLS] and [SEP] are the special tokens in BERT-based models (Devlin et al., 2019). [CLS] is used to represent the whole sentence, and [SEP] is used to separate different sentences, respectively. $h_1, \dots, h_{|h|}$ are the tokens of the head h , and $t_1, \dots, t_{|t|}$ are the tokenized tokens of the tail t . $[r]$ is registered as a new special token for a certain relation r . After feeding the serialized version of (h, r, t) into a BERT-based masked language model, the representation of the special token [CLS] is regarded as the representation of the whole triple. It is trained to distinguish positive triples with negative triples with cross entropy loss. Here x denotes a triple (h, r, t) , P_L models the distribution of the labeled dataset D_l , and θ is the set of parameters for KG-BERT. $P_\theta(y|x)$ denotes the probability after feeding the model prediction logits to softmax under parameter set θ . Then the optimization objective is as follows:

$$J(\theta) = \mathbb{E}_{x_l \sim P_L(x)}[-\log P_\theta(y|x_l)]. \quad (1)$$

4 Methods

In this section, we present the details of the framework of PseudoReasoner. A sketch illustration of the model is presented in Figure 2. To sum up, the procedure of PseudoReasoner can be summarized into the following steps:

- 1) Train a teacher model and a student model on the labeled dataset D_l (Section 4.1.1).
- 2) Use the teacher model to predict plausibility scores on triples from the unlabeled D_u . Triples with high/low plausibility scores within pre-defined intervals are given label 1/0 (Section 4.1.2).
- 3) Filter the pseudo labels with influence function with respect to the student model, and the student model’s predictions. (Section 4.1.3).
- 4) Finetune the student model on filtered pseudo labels from 3). (Section 4.2).

4.1 Pseudo Label Construction

4.1.1 Teacher Models

We use a teacher model pre-trained on the labeled dataset for labeling the unlabeled triples. We define plausibility scores of an unlabeled triple x as $\alpha(x)$, where the higher the score the more plausible the triple is regarded by the teacher model. We choose two different forms of teacher models as follows:

GPT2 (Radford et al., 2019): As negative sampling in the labeled dataset D_l is noisy, we aim to use an alternative model that avoids the negative part D_l^- . We finetune a (COMET) GPT2 language model, as the representative of generative family, on the positive part of the labeled dataset, D_l^+ , with a text generation task. For an (h, r, t) triple from D_l^+ , denote x as the serialized version of the triple, “ $h_1, \dots, h_{|h|}, [r], t_1, \dots, t_{|t|}$ ”. θ_{LM} denotes the trainable parameters in GPT2 language model (LM). We minimize the negative log likelihood of each triple as indicated in Equation (2):

$$L(x, \theta_{LM}) = -\frac{1}{|x|} \sum_{i=1}^{|x|} \log P(x_i|x_{<i}, \theta_{LM}). \quad (2)$$

Denote the optimized parameters as θ_{LM}^* . Here, the plausibility function $\alpha(x) = -L(x, \theta_{LM}^*)$, where the lower the loss, the higher the plausibility score by GPT2. Hence, for the triples from the unlabeled dataset, \mathcal{D}_u , we score every triple with Equation (2) on θ_{LM}^* .

KG-BERT: Besides GPT2, KG-BERT itself, a discriminative model, can be used as a teacher model. This teacher model learns θ^* from the labeled dataset D_l with cross entropy loss in Equation (1). For an instance $\{(h, r, t), y\} \in D_l$, denote $x = (h, r, t)$, we use $\alpha(x) = P_{\theta^*}(y=1|x)$ as x ’s plausibility score.

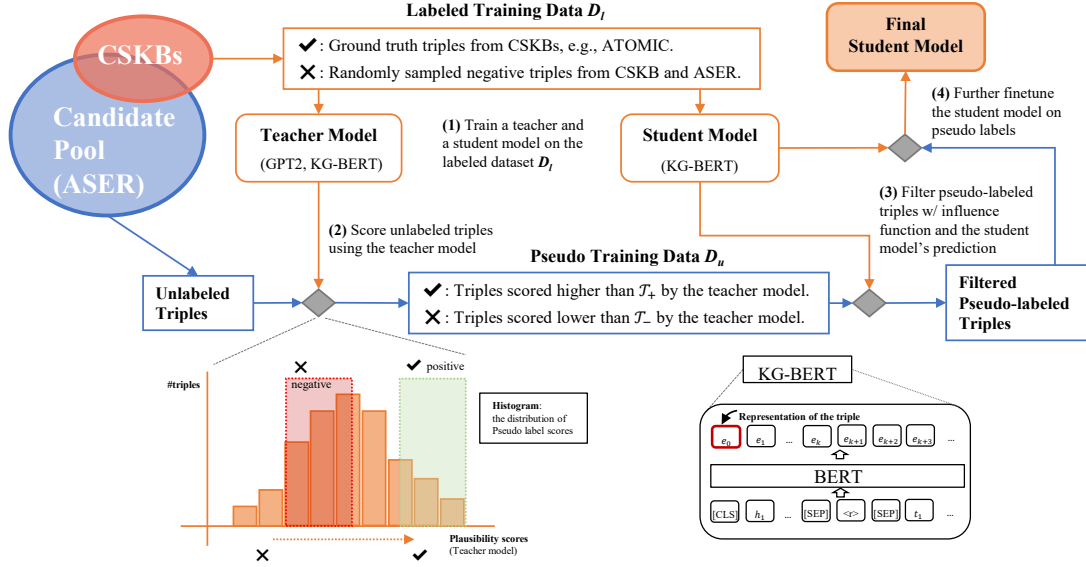


Figure 2: An end-to-end workflow of **PseudoReasoner**. Four steps in the figure are elaborated in Section 4.

4.1.2 Acquiring Pseudo Labels

The triples whose plausibility scores $\alpha(x)$ are between $[\mathcal{T}_{min}^-, \mathcal{T}_{max}^-]$ are labeled as negative, and the triples within $[\mathcal{T}_{min}^+, \mathcal{T}_{max}^+]$ are labeled as positive. Here $\mathcal{T}_{min}^- < \mathcal{T}_{max}^- < \mathcal{T}_{min}^+ < \mathcal{T}_{max}^+$. The reason that we introduce additional \mathcal{T}_{min}^- and \mathcal{T}_{max}^+ is that we want to filter out the triples that are treated over plausible or implausible by GPT2 to reduce potential *selection bias*. For example, GPT2 has been shown to provide low loss for repetitive patterns instead of the plausibility of the semantics (Brown et al., 2020). Details of the hyperparameter selection are in Appendix C.

4.1.3 Pseudo Label Filters

To further improve the quality of pseudo labels, we propose two filtering mechanisms on pseudo labels for better finetuning.

Influence Function. Filtering out detrimental training examples with influence function (Koh and Liang, 2017) can boost the model performance, as shown in Yang et al. (2020) and Han et al. (2020). A training example $z = ((h, r, t), y)$ will hurt the generalization ability of the model if including z in the training set results in a higher validation loss. Denote $L(\mathcal{Z}, \theta)$ as the loss function of dataset \mathcal{Z} under the parameter set θ . Then the loss under training set \mathcal{Z}_{train} is indicated in Equation (3):

$$L(\mathcal{Z}_{train}, \theta) = \frac{1}{|\mathcal{Z}_{train}|} \sum_{i=1}^{|\mathcal{Z}_{train}|} L(z_i, \theta). \quad (3)$$

Denote θ^* as the optimized parameters after training the model on \mathcal{Z}_{train} , and θ_{-z}^* as the optimized parameters after training the model on $\mathcal{Z}_{train} - \{z\}$.

Denote \mathcal{Z}_{val} as the validation set. The empirical criterion to determine z as a detrimental training example is Equation (4):

$$L(\mathcal{Z}_{val}, \theta^*) - L(\mathcal{Z}_{val}, \theta_{-z}^*) > 0. \quad (4)$$

The left-hand-side $L(\mathcal{Z}_{val}, \theta^*) - L(\mathcal{Z}_{val}, \theta_{-z}^*)$ can be approximated without retraining the model by influence function (Koh and Liang, 2017):

$$\mathcal{I}_{up,loss}(z) = -\nabla_{\theta} L(\mathcal{Z}_{val}, \theta^*)^{\top} H_{\theta^*}^{-1} \nabla_{\theta} L(z, \theta^*), \quad (5)$$

where $H_{\theta^*} = \frac{1}{|\mathcal{Z}_{train}|} \sum_{z_i \in \mathcal{Z}_{train}} \nabla_{\theta^*}^2 L(z_i, \theta^*)$ is the Hessian. We linearly approximate $\mathcal{I}_{up,loss}$ with inverse hessian-vector product (HVP) introduced in LiSSA (Agarwal et al., 2017) following (Koh and Liang, 2017). Details about influence function and the numerical approximation to calculate it can be found in the Appendix B. We filter out those examples with negative influence scores, which are harmful to the generalization of the model.

KG-BERT. As the student model we use is KG-BERT, when the pseudo labels from GPT2 are used, we can use the $P_{\theta^*}(y|x)$ produced from optimized KG-BERT as an additional filter to select pseudo labels. Specifically, pseudo labels $\{x = (h, r, t), y\}$ with $P_{\theta^*}(y|x) > 0.5$ are selected. This procedure can be viewed as ensembling GPT2 and KG-BERT.

4.2 PseudoReasoner Training

The objective function of KG-BERT on the labeled dataset is shown in Equation (1), and the objective function on pseudo labels can be written as:

$$J_U(\theta) = \mathbb{E}_{x_u \sim P_U(x)} \mathbb{E}_{\hat{y} \sim q(y|x_u)} [-\log P_{\theta}(\hat{y}|x_u)]. \quad (6)$$

Category	Model	all	In-domain	Out-of-domain (OOD)		Δ_{all}	Δ_{OOD}
			<i>Original Test Set</i>	<i>CSKB head + ASER tail</i>	<i>ASER edges</i>		
Supervised Learning	KG-BERT (BERT-base) <i>110M</i>	62.5	74.2	51.9	54.7	-	-
	KG-BERT (BERT-large) <i>340M</i>	67.7	74.5	58.7	62.0	-	-
	KG-BERT (DeBERTa-base) <i>100M</i>	64.5	73.2	54.0	57.0	-	-
	KG-BERT (DeBERTa-large) <i>350M</i>	69.2	77.6	59.9	61.8	-	-
	KG-BERT (BART-base) <i>139M</i>	65.1	74.7	54.7	56.6	-	-
	KG-BERT (BART-large) <i>406M</i>	70.4	<u>78.6</u>	62.8	64.2	-	-
	KG-BERT (RoBERTa-base) <i>110M</i>	68.0	76.3	59.3	59.8	-	-
	KG-BERT (RoBERTa-large) <i>340M</i>	<u>70.9</u>	78.0	<u>63.4</u>	<u>64.6</u>	-	-
	COMET (GPT2-small) <i>117M</i>	69.6	71.6	67.4	65.0	-	-
	COMET (GPT2-medium) <i>345M</i>	69.7	71.9	67.0	67.9	-	-
COMET (GPT2-large) <i>774M</i>	70.6	73.7	66.8	68.0	-	-	
COMET (GPT2-XL) <i>1558M</i>	70.7	74.6	66.7	67.6	-	-	
Semi-supervised Learning (RoBERTa-large)	UDA (TF-IDF)	71.7	78.0	65.1	65.9	+0.8	+1.5
	UDA (back-trans.)	71.6	78.6	64.2	66.2	+0.7	+1.2
	G-DAUG	71.7	78.5	64.8	65.5	+0.8	+1.2
	G-DAUG (COMET-distill)	72.2	78.6	65.9	66.9	+1.3	+2.4
	Noisy-student	72.4	79.3	65.3	66.7	+1.5	+2.0
Ours	PseudoReasoner (BERT-base)	67.9	76.0	56.1	64.2	+5.4	+6.9
	PseudoReasoner (RoBERTa-large)	74.2	80.1	69.5	69.3	+3.3	+5.3

Table 3: Results on the test set of the CSKB Population benchmark. For supervised learning baselines, we report the result of KG-BERT (Yao et al., 2019) with four backbone encoders and GPT2 (use LM loss to score triples). For semi-supervised learning (SSL) baselines, we study UDA (Xie et al., 2020a), G-DAUG (Yang et al., 2020), and Noisy-student (Xie et al., 2020b). The backbone encoders for SSL baselines are RoBERTa-large, which performs the best in the supervised setting. The number of parameters of backbone language models are presented as subscripts behind model names. Δ_{all} and Δ_{OOD} are the improvement on the “all” AUC and the Out-of-domain (OOD) AUC.

Here P_L and P_U are the distribution of the labeled and unlabeled dataset, respectively. $q(y|x)$ is the distribution of pseudo labels, modeled by the teacher model and filters. After finetuning KG-BERT initialized with θ^* on the filtered pseudo labels with Equation (6), we acquire θ^{*l} .

5 Experiments

5.1 Baselines

For the supervised learning setting, we use KG-BERT (Yao et al., 2019) and COMET (Bosselut et al., 2019) (GPT2) to perform CSKB Population. For KG-BERT, as it’s flexible to be adapted using different pre-trained encoders, we use BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), DeBERTa (He et al., 2021), and BART (Lewis et al., 2020) as the backbone language models. For BERT, RoBERTa, and DeBERTa, we use the embedding of the [CLS] token in KG-BERT as the representation of the whole triple. For BART, we follow the ways of doing sequence classification in the original paper (Lewis et al., 2020) to use the embedding of the end-of-sentence token in the decoder as the representation of the whole triple.

For the semi-supervised learning setting, we use the following baseline models:

Unsupervised Data Augmentation (UDA).

UDA (Xie et al., 2020a) uses consistency training to constrain the model to provide invariant predictions with noise added to the input. We adapt UDA into the framework of CSKB Population and uses TF-IDF word replacement and back-translation to provide noise to the input text to be fed into the consistency loss. More details are in the Appendix D.1.

Noisy Student.

Noisy student (Xie et al., 2020b) trains a student model with noise added during training iteratively. A teacher model is first trained to provide hard or soft pseudo labels for a student model to finetune together with the labeled dataset. Soft pseudo labels mean using logit scores after softmax as labels. Then the student model is iteratively re-used as the teacher model and a new student model is acquired through each iteration. Details can be found in the Appendix D.2.

Generative Data Augmentation (G-DAUG).

G-DAUG (Yang et al., 2020) leverages text generation language models to automatically generate pseudo training data examples for finetuning. Though

Teacher models	all	<i>Original Test Set</i>	<i>CSKB head + ASER tail</i>	<i>ASER edges</i>
N/A (Baseline)	70.9	78.0	63.4	64.6
w/o all filters				
RoBERTa-large	71.8	79.1	65.1	64.6
GPT2-small	72.3	79.5	65.1	65.9
GPT2-medium	72.6	79.5	65.8	68.0
GPT2-XL	72.8	80.1	66.3	66.0
w/ all filters				
RoBERTa-large	72.3	79.2	65.4	66.3
GPT2-small	73.7	79.8	67.5	68.8
GPT2-medium	74.1	81.3	68.1	69.0
GPT2-XL	74.2	80.1	69.5	69.3

Table 4: Performance of **PseudoReasoner** (RoBERTa-large) using different teacher models w/ or w/o filters.

it’s not designed for semi-supervised learning, we adapt it to our framework of pseudo labeling to serve as a semi-supervised learning baseline. We use COMET (GPT2-XL) finetuned on D_l to generate pseudo examples with heads from D_u . Then those pseudo labels are filtered with both influence function, diversity heuristics, and KG-BERT scores. Then those pseudo examples are used the same way as in our PseudoReasoner for further finetuning. We also try to replace COMET with COMET-distill (West et al., 2022), the COMET trained with distilled commonsense knowledge from GPT3, which has a better performance and capacity than the vanilla COMET. Details can be found in the Appendix D.3.

5.2 Experimental Settings

The learning rate for all models are set as $1e-5$, and the batch size is 64. We use the framework of Huggingface Transformers¹ to form our codebase. Early stopping is used where the best checkpoint is selected when the largest validation AUC is achieved. For all experiments, we report the average scores across three different random seeds.

For thresholding, we set the thresholds $\mathcal{T}_{min}^- = -4.0$, $\mathcal{T}_{max}^- = -3.7$, $\mathcal{T}_{min}^+ = -2.8$, $\mathcal{T}_{max}^+ = -2.0$, by roughly observing the data distribution and representative knowledge triples in different range of plausibility scores. We then randomly down-sample the pseudo examples such that the number is the same as the original training data. Ablations are provided in section 6.1.2

See more details about thresholding and hyperparameters of different models in Appendix C.

¹<https://huggingface.co/>

Filter	all	<i>Original Test Set</i>	<i>CSKB head + ASER tail</i>	<i>ASER edges</i>
baseline	70.9	78.0	63.4	64.6
w/o filter	72.8	80.1	66.3	66.0
+ influence	73.2	78.0	68.3	70.6
+ KG-BERT	73.7	79.5	68.9	68.6
+ both	74.2	80.1	69.5	69.3

Table 5: The effects of different filters on pseudo labels when KG-BERT (RoBERTa-large) is the backbone and GPT2-XL is the teacher model.

5.3 Results

The main results are shown in Table 3. We compare the results of both supervised learning and semi-supervised learning approaches with our proposed model PseudoReasoner. The “all” column presents the overall AUC across all testing examples and is the main metric of CSKB Population. We also separately present the AUC of different test set categories, In-domain (*Original Test Set*) and Out-of-domain (*CSKB head + ASER tail*, and *ASER edges*). In the last two columns, we report the increase by applying different semi-supervised learning approaches under the same backbone model, where Δ_{all} means the increase of “all” (AUC) metric and Δ_{OOD} means the increase of AUC for out-of-domain test sets.

For supervised-learning approaches, KG-BERT-based models mostly perform well on the In-domain test set while have poorer generalization ability to Out-of-domain test sets compared to COMET (GPT2). As GPT2 is only finetuned on the positive part of the dataset, it suffers less from the bias of negative sampling in D_l^- and has a better generalization on new knowledge. However, the drawback in In-domain reasoning hinders the overall performance of GPT2-XL from surpassing KG-BERT (RoBERTa-large), even with 4.5 times more parameters.

For semi-supervised learning baselines, all of them can increase the performance of the backbone KG-BERT (RoBERTa-large), especially for the out-of-domain split. However, the improvement on the In-domain part remains insignificant and the improvement on out-of-domain part is not as competitive as our PseudoReasoner. As we use the same code base and training method to train all semi-supervised learning methods, the main differences between PseudoReasoner and other SSL methods lie in the ways of processing the unlabeled dataset. We leave the detailed discussions in the next section (Section 6.2).

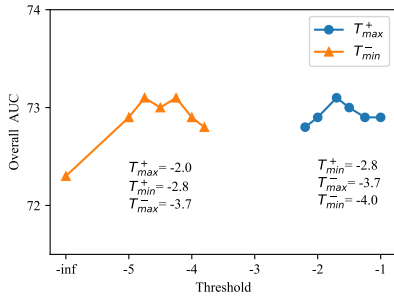


Figure 3: Ablation study on different \mathcal{T}_{max}^+ and \mathcal{T}_{min}^- .

6 Analysis and Discussions

In this section, we discuss the ablation study on model components, the comparisons with semi-supervised learning baselines, diversity analysis, and discussions about why PseudoReasoner works.

6.1 Ablation Study

We study the effects of different teacher models and the filters on pseudo labels.

6.1.1 Teacher Models

We compare four representative teacher models, KG-BERT (RoBERTa-large) and GPT2 (small, medium, and XL), on how pseudo labels provided by them can influence the model performance. The ablation results are shown in Table 4. From the comparison between KG-BERT (RoBERTa-large, 340M parameters) and GPT2-small (117M) and -medium (345M), which are of the same scale of model size, we find that GPT2 can perform consistently better than KG-BERT as teacher models. This can be validated by the out-of-distribution performance in Table 3 for the supervised learning baselines, where KG-BERT performs almost 3 points behind GPT2-medium in terms of OOD AUC. This ablation indicates the importance of powerful generalizable teacher models on pseudo labeling.

6.1.2 Thresholding

We study the sensitivity of thresholds in Figure 3. In this ablation, for simplicity, we set different \mathcal{T}_{max}^+ for positive pseudo examples, and sample the same amount of triples as the original training set whose $\alpha(x) < \mathcal{T}_{max}^+$ in descending order. We do the same ablation study on \mathcal{T}_{min}^- for negative pseudo examples. When tuning one threshold, other thresholds are fixed as in section 5.2. The pseudo labels under different thresholds are directly used for PseudoReasoner without filtering, and we plot the test set AUC given different thresholds.

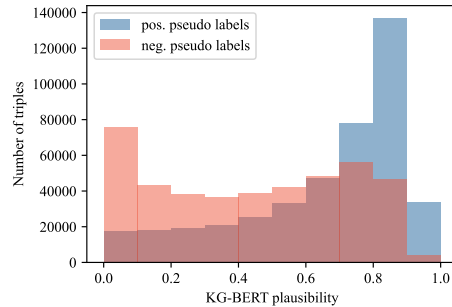


Figure 4: KG-BERT plausibility distribution for positive/negative pseudo labels provided by GPT2.

	w/o filter	influence	KG-BERT	both
UDA	3.3 M	-	-	-
G-DAUG	1.1M	399.8K	323.3K	160.1K
PseudoReasoner	932.6K	408.5K	373.0K	170.1K
Original	1.1 M	-	-	-

Table 6: Number of pseudo examples used in experiments for semi-supervised methods. The “Original” row indicates the number of training examples in the original training set.

We see that the resulting AUC is stable within certain ranges of \mathcal{T}_{max}^+ and \mathcal{T}_{min}^- . While when we set \mathcal{T}_{min}^- to $-\infty$, which indicates no thresholds for negative examples are set, the performance drops drastically.

6.1.3 Pseudo Label Filter

We conduct experiments with different combinations of filtering mechanisms in Table 5 for KG-BERT (RoBERTa-large). We can see that both filters (influence function and KG-BERT probability) benefits the model performance, while KG-BERT probability contributes to a more substantial improvement. Figure 4 shows an illustration of the KG-BERT plausibility $\alpha(x) = P_{\theta}^*(y=1|x)$ of positive/negative pseudo examples provided by GPT2. The positive pseudo examples tend to be scored higher by KG-BERT than negative pseudo examples. Adding KG-BERT probability as an additional filter with the labels provided by GPT2 is similar to an ensembling procedure.

6.2 Comparisons with Other Semi-supervised Learning Methods

6.2.1 Computational Cost

The number of pseudo examples used for semi-supervised-learning baselines are listed in Table 6. We basically use the same scale of unfiltered pseudo examples for G-DAUG and PseudoReasoner, while use 3 times more unlabeled examples for UDA as it requires more unlabeled data. Specifically, G-DAUG (COMET-distill) leverages

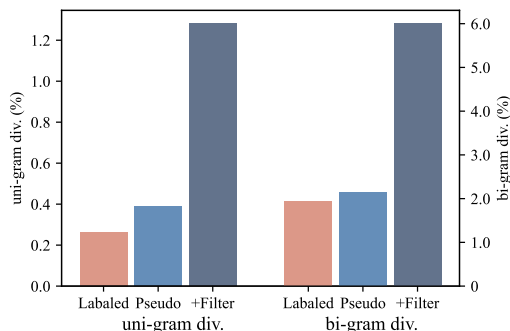


Figure 5: Diversity analysis with the proportion of unique uni/bi-grams in the labeled dataset, the pseudo labels, and the filtered pseudo labels. With filtering, the diversity can be significantly improved.

the distilled knowledge from GPT3 (West et al., 2022), which further magnify the computational cost to more orders of magnitude. In all, under the same scale of pseudo labels, PseudoReasoner can achieve far better results than UDA and G-DAUG.

6.2.2 Analysis

In UDA, though robustness can be improved with consistency loss on noised inputs, there are no *new* commonsense knowledge added to the training procedure, making it hard for the model to be equipped with novel knowledge reasoning ability.

For G-DAUG, the critical part lies in the generation of negative examples. We finetune two separate GPT2 on D_l^+ and D_l^- , and the one finetuned on D_l^- is used to generate negative examples. Compared with the GPT2 finetuned on D_l^+ , the GPT2 finetuned on D_l^- is of a relatively lower quality as triples in D_l^- don't follow specific commonsense patterns. We check the text generation quality of GPT2-XL finetuned on D_l^+ and D_l^- , and find that the BLEU-2 scores on two corresponding held-out test sets are 0.23 and 0.10, indicating the generation of negative examples are of a lower quality.

For Noisy-student, the main differences between our PseudoReasoner is that they use KG-BERT to provide pseudo labels, and they are soft pseudo labels. The main reason behind is that as soft labels are used, the teacher model has to be a discriminative model such as KG-BERT, which has a poor generalization ability than GPT2 used in our PseudoReasoner. Moreover, similar to the case in UDA, noisy is not a dominate factor in CSKB Population, while more high-quality novel commonsense knowledge matters more.

6.3 Semantic Diversity Analysis

An important contribution of PseudoReasoner is that we extend the knowledge space for training

from limited CSKBs to a more broad unlabeled resource. We use the proportion of unique uni-grams and bi-grams as an indicator of semantic diversity to measure the scale to which models are exposed to diverse novel knowledge. Figure 5 shows that after filtering with influence function and KG-BERT probability, the diversity can be improved by around 3 to 4 times than the labeled dataset.

6.4 Relationship with Knowledge Distillation

While knowledge distillation focuses on distilling knowledge from larger models to smaller ones, our method does not necessarily need the teacher model to be larger. The teacher GPT2-medium, which is of the same size as the student KG-BERT, can work pretty well and is comparable to GPT2-XL.

7 Conclusion

In this paper, we propose a semi-supervised learning framework for CSKB Population based on pseudo labels. Using a teacher model and a special filtering mechanism on pseudo labels, we achieve the state-of-the-art of CSKB Population in terms of both in-domain and out-of-domain performance. Experiments also show that our CSKB Population benefits more from high-quality novel knowledge than other semi-supervised learning techniques such as noise and consistency training. This work brings a new perspective of improving out-of-domain generalizable commonsense reasoning ability on CSKBs.

Acknowledgement

The authors of this paper were supported by the NSFC Fund (U20B2053) from the NSFC of China, the RIF (R6020-19 and R6021-20) and the GRF (16211520) from RGC of Hong Kong, the MHKJFS (MHP/001/19) from ITC of Hong Kong and the National Key R&D Program of China (2019YFE0198200) with special thanks to HKMAAC and CUSBLT, and the Jiangsu Province Science and Technology Collaboration Fund (BZ2021065). We thank the support from the UGC Research Matching Grants (RMGS20EG01-D, RMGS20CR11, RMGS20CR12, RMGS20EG19, RMGS20EG21). We also thank NVIDIA AI Technology Center (NVAITC) for the support of computational resources for conducting the experiments.

Limitations

The main limitation is that for the labeling procedure of the teacher model, we still need a hard threshold \mathcal{T} to determine whether the label is 0 or 1. This would limit the method from generalize across different tasks, where different thresholds should be tuned separately. One solution and future work would be automatically learn the thresholding, or improve the teacher-student interaction by training them jointly in an end-to-end manner such as Meta Pseudo Label (Pham et al., 2021).

References

- Naman Agarwal, Brian Bullins, and Elad Hazan. 2017. Second-order stochastic optimization for machine learning in linear time. *J. Mach. Learn. Res.*, 18:116:1–116:40.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Antoine Bosselut, Ronan Le Bras, and Yejin Choi. 2021. Dynamic neuro-symbolic knowledge graph construction for zero-shot commonsense question answering. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 4923–4931. AAAI Press.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. COMET: commonsense transformers for automatic knowledge graph construction. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4762–4779. Association for Computational Linguistics.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Junkun Chen, Renjie Zheng, Atsuhito Kita, Mingbo Ma, and Liang Huang. 2021. Improving simultaneous translation by incorporating pseudo-references with fewer reorderings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 5857–5864. Association for Computational Linguistics.
- R. Dennis Cook and Sanford Weisberg. 1982. *Residuals and influence in regression*. Monogr. Stat. Appl. Probab. CRC Press, Boca Raton, FL.
- Ernest Davis. 1990. *Representations of commonsense knowledge*. notThenot Morgan Kaufmann series in representation and reasoning. Morgan Kaufmann.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pages 248–255. IEEE Computer Society.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Tianqing Fang, Weiqi Wang, Sehyun Choi, Shibo Hao, Hongming Zhang, Yangqiu Song, and Bin He. 2021a. Benchmarking commonsense knowledge base population with an effective evaluation dataset. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 8949–8964. Association for Computational Linguistics.
- Tianqing Fang, Hongming Zhang, Weiqi Wang, Yangqiu Song, and Bin He. 2021b. DISCOS: bridging the gap between discourse knowledge and commonsense knowledge. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, pages 2648–2659. ACM / IW3C2.
- Suchin Gururangan, Tam Dang, Dallas Card, and Noah A. Smith. 2019. Variational pretraining for semi-supervised text classification. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5880–5894. Association for Computational Linguistics.

- Xiaochuang Han, Byron C. Wallace, and Yulia Tsvetkov. 2020. Explaining black box predictions and unveiling data artifacts through influence functions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5553–5563. Association for Computational Linguistics.
- Junxian He, Jiatao Gu, Jiajun Shen, and Marc’Aurelio Ranzato. 2020. Revisiting self-training for neural sequence generation. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Mutian He, Tianqing Fang, Weiqi Wang, and Yangqiu Song. 2022. Acquiring and modelling abstract commonsense knowledge via conceptualization. *arXiv preprint arXiv:2206.01532*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. Deberta: decoding-enhanced bert with disentangled attention. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Kaiyu Huang, Junpeng Liu, Degen Huang, Deyi Xiong, Zhuang Liu, and Jinsong Su. 2021. Enhancing chinese word segmentation via pseudo labels for practicability. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 4369–4381. Association for Computational Linguistics.
- Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. (comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 6384–6392. AAAI Press.
- Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. 2019. Label propagation for deep semi-supervised learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 5070–5079. Computer Vision Foundation / IEEE.
- Stanislaw Jastrzebski, Dzmitry Bahdanau, Seydarian Hosseini, Michael Noukhovitch, Yoshua Bengio, and Jackie Cheung. 2018. Commonsense mining as knowledge base completion? a study on the impact of novelty. In *Proceedings of the Workshop on Generalization in the Age of Deep Learning*, pages 8–16, New Orleans, Louisiana. Association for Computational Linguistics.
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1885–1894. PMLR.
- Dong-Hyun Lee et al. 2013. 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.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.
- Xiang Li, Aynaz Taheri, Lifu Tu, and Kevin Gimpel. 2016. Commonsense knowledge base completion. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Hugo Liu and Push Singh. 2004. Conceptnet—a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Chaitanya Malaviya, Chandra Bhagavatula, Antoine Bosselut, and Yejin Choi. 2020. Commonsense knowledge base completion with structural and semantic context. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 2925–2933. AAAI Press.
- Nasrin Mostafazadeh, Aditya Kalyanpur, Lori Moon, David Buchanan, Lauren Berkowitz, Or Biran, and Jennifer Chu-Carroll. 2020. Glucose: Generalized and contextualized story explanations. In *The Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V. Le. 2021. Meta pseudo labels. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 11557–11568. Computer Vision Foundation / IEEE.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners.

- Itsumi Saito, Kyosuke Nishida, Hisako Asano, and Junji Tomita. 2018. Commonsense knowledge base completion and generation. In *Proceedings of the 22nd Conference on Computational Natural Language Learning, CoNLL 2018, Brussels, Belgium, October 31 - November 1, 2018*, pages 141–150. Association for Computational Linguistics.
- Maarten Sap, Ronan Le Bras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3027–3035.
- Weiwei Shi, Yihong Gong, Chris Ding, Zhiheng Ma, Xiaoyu Tao, and Nanning Zheng. 2018. Transductive semi-supervised deep learning using min-max features. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part V*, volume 11209 of *Lecture Notes in Computer Science*, pages 311–327. Springer.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Unsupervised commonsense question answering with self-talk. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4615–4629. Association for Computational Linguistics.
- Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2071–2080. JMLR.org.
- Bin Wang, Guangtao Wang, Jing Huang, Jiakuan You, Jure Leskovec, and C.-C. Jay Kuo. 2021. Inductive learning on commonsense knowledge graph completion. In *International Joint Conference on Neural Networks, IJCNN 2021, Shenzhen, China, July 18-22, 2021*, pages 1–8. IEEE.
- Kexin Wang, Nandan Thakur, Nils Reimers, and Iryna Gurevych. 2022. GPL: generative pseudo labeling for unsupervised domain adaptation of dense retrieval. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena D. Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. Symbolic knowledge distillation: from general language models to commonsense models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 4602–4625. Association for Computational Linguistics.
- Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. 2020a. Unsupervised data augmentation for consistency training. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Qizhe Xie, Minh-Thang Luong, Eduard H. Hovy, and Quoc V. Le. 2020b. Self-training with noisy student improves imagenet classification. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10684–10695. Computer Vision Foundation / IEEE.
- I. Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. 2019. Billion-scale semi-supervised learning for image classification. *CoRR*, abs/1905.00546.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. G-daug: Generative data augmentation for commonsense reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 1008–1025. Association for Computational Linguistics.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. KG-BERT: BERT for knowledge graph completion. *CoRR*, abs/1909.03193.
- Hongming Zhang, Xin Liu, Haojie Pan, Haowen Ke, Jiefu Ou, Tianqing Fang, and Yangqiu Song. 2022. Aser: Towards large-scale commonsense knowledge acquisition via higher-order selectional preference over eventualities. *Artificial Intelligence*, 309:103740.
- Hongming Zhang, Xin Liu, Haojie Pan, Yangqiu Song, and Cane Wing-Ki Leung. 2020. ASER: A large-scale eventuality knowledge graph. In *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*, pages 201–211. ACM / IW3C2.
- Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schölkopf. 2003. Learning with local and global consistency. In *Advances in Neural Information Processing Systems 16 [Neural Information Processing Systems, NIPS 2003, December 8-13, 2003, Vancouver and Whistler, British Columbia, Canada]*, pages 321–328. MIT Press.

A Additional Details on the CSKB Population Benchmark

Commonsense Knowledge Base (CSKB) Population aims at automatically populating commonsense knowledge defined in source CSKBs on an unlabeled candidate eventuality knowledge graph, ASER. We provide additional details about the selection of CSKBs and the eventuality knowledge graph, and the details about the evaluation of CSKB Population in this section.

A.1 CSKB

CSKB Population studies commonsense relations among general events. To form a general aligned CSKB, Fang et al. (2021a) selected ConceptNet (Liu and Singh, 2004) (the event-related relations are selected), ATOMIC (Sap et al., 2019), ATOMIC₂₀ (Hwang et al., 2021) (the newly added relations beyond ATOMIC are selected), and GLUCOSE (Mostafazadeh et al., 2020). These four sources of CSKBs are in the form of free-text and are structured as triplet forms $((h, r, t))$. Normalization processes are conducted to align the four CSKBs together. The overall structure of the aligned CSKB is based on the formats in ATOMIC, where the events are person-centric sentences with *PersonX* and *PersonY* as subjects. In ConceptNet, *PersonX* is prepended on the *predicate-object* pairs to make them complete sentences. For example, a triple in ConceptNet (lie, HasSubEvent, make up story) is converted to (PersonX lies, HasSubEvent, PersonX makes up story). The *SomeoneA* and *SomeoneB* are converted to *PersonX* and *PersonY* accordingly in GLUCOSE. The relationships in GLUCOSE are converted to ATOMIC formats according to the official conversion rules defined in Table 7 in Mostafazadeh et al. (2020).

A summary of the relations studied in the aligned CSKBs is shown in Table 7.

A.2 ASER

The candidate knowledge graph for populating commonsense knowledge, ASER (Zhang et al., 2020, 2022), is a large-scale eventuality-centric knowledge graph that provides explicit discourse relationships between eventualities. The core part of ASER is used, where 10M discourse edges among 27M eventualities are included. An example of an ASER edge is: (“I am hungry,” *Result*, “I have lunch”), and such an discourse edge can be a potential *xWant* relation in CSKBs. Such

Relation	ATOMIC ₂₀	ConceptNet	GLUCOSE
<i>oEffect</i>	21,497	0	7,595
<i>xEffect</i>	61,021	0	30,596
<i>gEffect</i>	0	0	8,577
<i>oWant</i>	35,477	0	1,766
<i>xWant</i>	83,776	0	11,439
<i>gWant</i>	0	0	5,138
<i>oReact</i>	21,110	0	3,077
<i>xReact</i>	50,535	0	13,203
<i>gReact</i>	0	0	2,683
<i>xAttr</i>	89,337	0	7,664
<i>xNeed</i>	61,487	0	0
<i>xIntent</i>	29,034	0	8,292
<i>isBefore</i>	18,798	0	0
<i>isAfter</i>	18,600	0	0
<i>HinderedBy</i>	87,580	0	0
<i>xReason</i>	189	0	0
<i>Causes</i>	0	42	26,746
<i>HasSubEvent</i>	0	9,934	0
Total	578,252	10,165	126,776

Table 7: Relation distribution statistics for different CSKBs. The table is the same as in the Table 4 of Fang et al. (2021a).

detailed conversion rules from discourse relations to commonsense relations are provided in Fang et al. (2021a). Such conversion rules make the total number of unlabeled candidate edges around 200M.

A.3 Evaluation Set

For the ground truth commonsense triples from the CSKBs, they are split in to train, development, and test sets using the original split by their own papers. For example, in ATOMIC, the test set split is based on the skeleton words of the heads such that similar heads are split into the same dataset. The test set here is denoted as the *original test set*.

The final evaluation set comprises of three parts, one sampled from the original automatically constructed test set as above (denoted as “*Original Test Set*”), one sample from the edges where heads are from CSKBs and tails are from ASER (denoted as “*CSKB head + ASER tail*”), and one from ASER solely (denoted as “*ASER edges*”). The total number of final evaluation set is around 32K, and they are manually annotated using Amazon Mechanical Turk.

B Influence Function

Filtering out detrimental training examples with influence function (Koh and Liang, 2017) can boost the model performance, as shown in Yang et al. (2020) and Han et al. (2020). A training example

$z = ((h, r, t), y)$ will hurt the generalization ability of the model if including z in the training set results in a higher validation loss. Denote $L(\mathcal{Z}, \theta)$ as the loss function of dataset \mathcal{Z} under the parameter set θ . Then the loss under the training set \mathcal{Z}_{train} of cardinality n is indicated in Equation (3) in the main body (also shown below).

$$L(\mathcal{Z}_{train}, \theta) = \frac{1}{n} \sum_{i=1}^n L(z_i, \theta).$$

Denote θ^* as the optimized parameters after training the model on \mathcal{Z}_{train} , and θ_{-z}^* as the optimized parameters after training the model on $\mathcal{Z}_{train} - \{z\}$. Denote \mathcal{Z}_{val} as the validation set. The empirical criterion to determine z as a detrimental training example is Equation (4) in the main body (also shown below).

$$L(\mathcal{Z}_{val}, \theta^*) - L(\mathcal{Z}_{val}, \theta_{-z}^*) > 0.$$

Influence function gives an efficient approximation for the above quantity, whose idea is to compute the change of the optimized parameters thus the validation loss if z were upweighted by a small ϵ . It introduces new parameters $\theta_{\epsilon, z}^* = \arg \min_{\theta} \frac{1}{n} \sum_{i=1}^n L(z_i, \theta) + \epsilon L(z, \theta)$. Cook and Weisberg (1982) gives us the influence of upweighting z on the optimized parameters θ^* :

$$\mathcal{I}_{up, params}(z) = \left. \frac{d\theta_{\epsilon, z}^*}{d\epsilon} \right|_{\epsilon=0} = -H_{\theta^*}^{-1} \nabla_{\theta} L(z, \theta^*)$$

where $H_{\theta^*} = \frac{1}{n} \sum_{i=1}^n \nabla_{\theta}^2 L(z_i, \theta^*)$ is the Hessian matrix. Thus, applying the chain rule, we calculate the influence of upweighting z on the validation loss:

$$\begin{aligned} \mathcal{I}_{up, loss}(z) &= \left. \frac{dL(\mathcal{Z}_{val}, \theta_{\epsilon, z}^*)}{d\epsilon} \right|_{\epsilon=0} \\ &= \nabla_{\theta} L(\mathcal{Z}_{val}, \theta^*)^{\top} \left. \frac{d\theta_{\epsilon, z}^*}{d\epsilon} \right|_{\epsilon=0} \\ &= -\nabla_{\theta} L(\mathcal{Z}_{val}, \theta^*)^{\top} H_{\theta^*}^{-1} \nabla_{\theta} L(z, \theta^*). \end{aligned}$$

Since removing z is equivalent to unweighting it by $\epsilon = -\frac{1}{n}$, small when n is large, we can linearly approximate $L(\mathcal{Z}_{val}, \theta^*) - L(\mathcal{Z}_{val}, \theta_{-z}^*)$ in Equation (4) as $-\epsilon \mathcal{I}_{up, loss}(z)$ ($= \frac{1}{n} \mathcal{I}_{up, loss}(z)$).

In practice, we also linearly approximate $\mathcal{I}_{up, loss}$ with inverse Hessian-vector product (HVP) introduced in LiSSA (Agarwal et al., 2017) following (Koh and Liang, 2017).

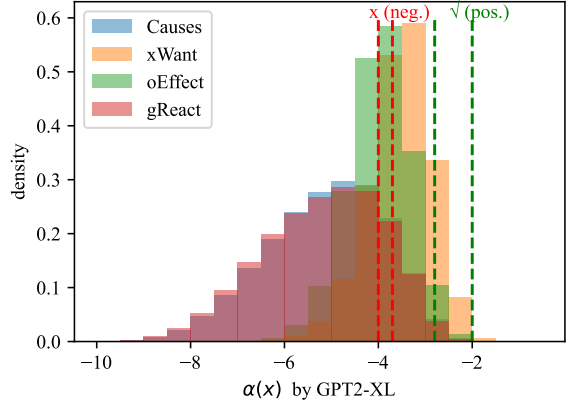


Figure 6: GPT2-XL plausibility distribution ($-\alpha(x)$) of unlabeled data with relation Causes, xWant, oEffect, and gReact. The triples within the red dashed lines are considered negative pseudo examples, and the triples within the green dashed lines are considered as positive pseudo examples.

C Additional Details on Hyper-parameters

We present additional details about hyperparameters in the thresholding of the GPT2-XL teacher model and the details of the backbone models.

C.1 Thresholding in Pseudo Labels

We present the GPT2-XL plausibility scores on the unlabeled D_u in Figure 6 for four representative commonsense relations, Causes, xWant, oEffect, and gReact. We set the thresholds $\mathcal{T}_{min}^- = -4.0$, $\mathcal{T}_{max}^- = -3.7$, $\mathcal{T}_{min}^+ = -2.8$, $\mathcal{T}_{max}^+ = -2.0$, by roughly observing the data distribution and representative triples in different range of plausibility scores. The filtered triples are further sampled such that the number of pseudo examples per relation equals the number of the original examples for the corresponding relations. The pseudo examples that we use for experiments can be downloaded at our GitHub repository.

C.2 Model Training Details

For KG-BERT, when BERT, RoBERTa, and DeBERTa is used, an (h, r, t) triple is converted to “[CLS], $h_1, \dots, h_{|h|}$, [SEP], $[r]$, [SEP], $t_1, \dots, t_{|t|}$ ”. Here, [CLS] and [SEP] are the special tokens in BERT-based models that are used to represent the whole sentence, and the special token to separate different sentences, respectively. $h_1, \dots, h_{|h|}$ is the tokenized tokens of the head h , and $t_1, \dots, t_{|t|}$ is the tokenized tokens of the tail t . $[r]$ is registered as a new special token for a certain relation r .

The embedding of [CLS] in the last hidden layer of the BERT-based encoder is used as the final representation of the whole triple.

When BART is used as the backbone, an (h, r, t) triple is converted to “<s>, $h_1, \dots, h_{|h|}$, $[r]$, $t_1, \dots, t_{|t|}$, </s>”, where <s> and </s> are the special tokens in BART to represent starting-of-sentence and end-of-sentence. The sentence is fed into the encoder of BART and also the decoder of BART. The embedding of the end-of-sentence token </s> in the last hidden layer of the decoder is used as the embedding of the whole triple. This method is the default usage as in the original paper of BART for text classification.

For GPT2, we adopt the text generation setting in COMET (Bosselut et al., 2019), where the task is defined as give the head and relation to generate tail. Specifically, an (h, r, t) triple is converted to “[SOS], $h_1, \dots, h_{|h|}$, $[r]$, $t_1, \dots, t_{|t|}$, [EOS]”, where [SOS] and [EOS] are special tokens representing start-of-sentence and end-of-sentence. The negative GPT2 loss of the whole sentence (Equation (2)) is used as the plausibility of the triple. Such a GPT2 can also be used to generate positive commonsense tails in the G-DAUG baseline. We adapt the codes² from ATOMIC₂₀²⁰ to our code base.

We build all models based on PyTorch and all transformer models from Huggingface transformers. For training, we use 1e-5 as the default learning rate for all training stages. The batch size is set to 64 for all experiments. The maximum sequence length is set to 30, which can cover over 99% of the serialized triples. We use ADAM as the optimizer to train all models. We evaluate the validation AUC every 250 steps and store the checkpoint when the largest validation AUC is achieved. We run all experiments with three random seeds (100, 101, and 102) and report the average. The KG-BERT related experiments are run on NVIDIA RTX 3080 Ti, and the GPT2 related experiments are run on NVIDIA A40.

D Details of Baselines

We present more details for the training of semi-supervised learning baselines in this section.

D.1 UDA

In the original paper of UDA, two data augmentation methods, back-translation and TF-IDF replace-

ment, are introduced for natural language tasks. We use both augmentation methods and compare the results. For each augmentation methods, original serialized unlabeled candidates x are duplicated and augmented, then both original and augmented candidates are used for consistency training.

For back-translation, we employ OPUS translation models³ between English and French. Specifically, in each candidate, the “PersonX” is replaced with Alice, “PersonY” with Charles, and “PersonZ” with Francisco, then candidates are forward-translated to French and after then backward-translated to English. Selected names Alice, Charles, and Francisco are common first names shared by French and English and remain unchanged after back-translation, thus it ensures that the personal pronouns PersonX, PersonY, and PersonZ in unlabeled candidates can be fully recovered after back-translation. Also, we observe that almost newly registered relation tokens survive after the back-translation process, bringing the opportunity for KG-BERT to get familiar to these tokens in consistency training. For TF-IDF replacement, we set the replacement probability as 0.1 to avoid cracking the main semantics of the original sentences.

We present more hyperparameter tuning here for UDA. The objective function of UDA is as follows in Equation (7), where p_L and p_U are the distribution of the labeled and unlabeled dataset, y_1 is the corresponding label of x_1 , p_θ is the predicted probability by the backbone model under parameter θ , $q(\hat{x}|x_2)$ is a data augmentation transformation, $\hat{\theta}$ is a fixed copy of the current parameters θ , and CE indicates cross entropy. λ is used to control the weight of unsupervised consistency loss. Besides λ , in the training process, there is another parameter r indicating the ratio of unsupervised examples in a batch. We conduct several experiments with grid search to find the best combinations of λ and r as shown in Table 8 when using TF-IDF replacement as the augmentation method. We report the results of the best combinations in the main text of the paper.

$$J(\theta) = \mathbb{E}_{x_1 \sim p_L(x)} [-\log p_\theta(y_1|x_1)] + \lambda \mathbb{E}_{x_2 \sim p_U(x)} \mathbb{E}_{\hat{x} \sim q(\hat{x}|x_2)} [\text{CE}(p_{\hat{\theta}}(y|x_2) || p_\theta(y|\hat{x}))] \quad (7)$$

The original implementation of UDA is based on Tensorflow⁴. We re-implement their code with

²<https://github.com/allenai/comet-atomic-2020>

³<https://huggingface.co/Helsinki-NLP>

⁴<https://github.com/google-research/uda>

$r \backslash \lambda$	0.3	1	3
0.3	71.0	71.4	71.2
1	71.4	71.7	71.2
3	71.2	71.2	71.4

Table 8: Overall AUC score with respect to different pairs of unsupervised example ratio r and consistency loss coefficient λ .

pytorch and aggregate it into our code base.

D.2 Noisy Student

For Noisy Student, they use a teacher model finetuned on the labeled dataset first to provide initial pseudo labels. Then student models are finetuned with both the labeled dataset the pseudo labels iteratively when the latest pseudo labels are provided by the latest student model. Specifically, noise is added to the input during the stage for student models’ training.

In this section, we provide some detailed ablation studies on the setting of Noisy Student. First, we study the role of noise in Noisy Student in Table 9 when noise is added or removed to the output logits of the model at the first training iteration. Results show that there can be a 0.3 points improvement on the overall AUC when dropout is added with probability either 0.1 or 0.5, showing that noise can indeed force the models to learn harder. Second, we study the performance across different iterations of Noisy Student. In table 9, the performance reaches the peak at iteration 3, indicating a balance of the student model’s learning ability and generalizable pseudo labeling ability. We report the best result in the result table of the main body.

D.3 G-DAUG

For Generative Data Augmentation, we check the performance under different sizes of COMET (use GPT2-small and GPT2-XL as representatives) in Table 10. Also, we check the three filtering mechanisms, influence function, diversity, and KG-BERT probability. The diversity filter is a heuristic algorithm that includes diverse training examples one by one as defined in the original paper. We find that still the KG-BERT filter is the most effective filter and can boost the performance more than the other two. Also, G-DAUG also benefits more from a more powerful text generation model (GPT2-XL).

	all	<i>Original Test Set</i>	<i>CSKB head + ASER tail</i>	<i>ASER edges</i>
KG-BERT (RoBERTa-large)	70.9	78.0	63.4	64.6
1st Iteration				
- w/o noisy	71.8	78.5	64.4	65.8
- dropout ($p=0.1$)	72.1	78.8	64.8	66.6
- dropout ($p=0.5$)	72.1	79.3	64.7	66.2
Num of Iterations				
1	72.1	79.3	64.7	66.2
2	72.1	78.6	64.6	66.8
3	72.4	79.3	65.3	66.7
4	72.3	79.4	65.1	66.4

Table 9: Ablation study on the effect of noisy in the Noisy Student model. p means the dropout probability.

	all	<i>Original Test Set</i>	<i>CSKB head + ASER tail</i>	<i>ASER edges</i>
KG-BERT (RoBERTa-large)	70.9	78.0	63.4	64.6
G-DAUG (GPT2-small)				
no filter	71.3	78.6	63.7	65.3
+ influence	70.6	76.8	63.0	66.5
+ diversity	71.4	78.5	64.2	65.3
+ KG-BERT	71.5	78.7	64.4	64.9
G-DAUG (GPT2-XL)				
no filter	71.5	78.7	64.2	65.4
+ influence	71.3	78.4	64.4	65.2
+ diversity	71.7	78.7	64.7	65.5
+ KG-BERT	71.7	78.5	64.8	65.5

Table 10: Effects of different filtering methods of the G-DAUG (Yang et al., 2020) baseline. We use RoBERTa-large as the base language model for KG-BERT. For example, “+ KG-BERT” means filter the augmented dataset with KG-BERT loss as illustrated in Section 4.1.3.