PRE-TRAINING TEXT-TO-TEXT TRANSFORMERS TO WRITE AND REASON WITH CONCEPTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Pretrained language models (PTLM) have achieved impressive results in a range of natural language understanding (NLU) and generation (NLG) tasks that require a syntactic and semantic understanding of the text. However, current pre-training objectives such as masked token prediction (for BERT-style PTLMs) and masked span infilling (for T5-style PTLMs) do not explicitly model the relational and compositional commonsense knowledge about everyday concepts, which is crucial to many downstream tasks requiring commonsense reasoning. To augment PTLMs with common sense, we propose generative and contrastive objectives as intermediate self-supervised pre-training tasks between general pre-training and downstream task-specific fine-tuning. We also propose a joint training framework to unify generative and contrastive objectives so that these objectives can be more effective. Our proposed objectives can pack more commonsense knowledge into the parameters of a pre-trained text-to-text transformer without relying on external knowledge bases, yielding better performance on both NLU and NLG tasks. We apply our method on a pre-trained T5 model in an intermediate task transfer learning fashion to train a *concept-aware language model* (**CALM**)[1], and experiment with five commonsense benchmarks (four NLU tasks and one NLG task). Experimental results show that CALM outperforms baseline methods by a consistent margin.

1 INTRODUCTION

Pre-trained language models (PLTMs) such as BERT (Devlin et al., 2018) and T5 (Raffel et al., 2019) have revolutionized the field of NLP, yielding impressive performance on various conventional natural language understanding (NLU) and generation (NLG) tasks. BERT and its novel variants such as RoBERTa (Liu et al., 2019) and ALBERT (Lan et al., 2019) capture syntactical and semantic knowledge mainly from the pre-training task of *masked language modeling*, while T5-style models such as BART (Lewis et al., 2019) instead focus on *masked span infilling* tasks. Though yielding better performance on downstream tasks, these pre-training objectives, however, do not explicitly guide the models to reason with concept-centric commonsense knowledge from language. This leaves room for equipping current PTLMs with richer commonsense reasoning ability.

For example, consider a multi-choice question "*What do you fill with ink to write notes on a piece of copy paper? (A) fountain pen (B) pencil case (C) printer (D) notepad*". The current state-of-the-art question answering model, UnifiedQA (Khashabi et al., 2020), which was fine-tuned on T5-large with multiple datasets, still predicts '(C) *printer*' as its answer. The model may be overly sensitive to the *co-occurrence* between phrases in question sentence like '*ink*' and '*copy paper*' and the answer choice '*printer*', but fails to reason with the concept-centric knowledge that '*fountain pen*' is a writing instrument that needs to be filled with '*ink*'. Such mistake in commonsense reasoning becomes a bottleneck for current PTLMs (Davis & Marcus, 2015). Towards augmenting PTLMs with more knowledge, prior works mainly focus on training larger models (Brown et al., 2020), adding knowledge-specific architectures (Peters et al., 2019), or incorporating knowledge bases for pre-training (Xiong et al., 2020). In this paper, we instead look to explicitly *teach* pre-trained models to write and reason with common concepts through novel pre-training strategies.

[1]Code and data have been uploaded and will be published: `https://anonymous.4open.science/repository/6fdeed55-ec2c-4ffa-aee8-0cc3b7f5ade5`

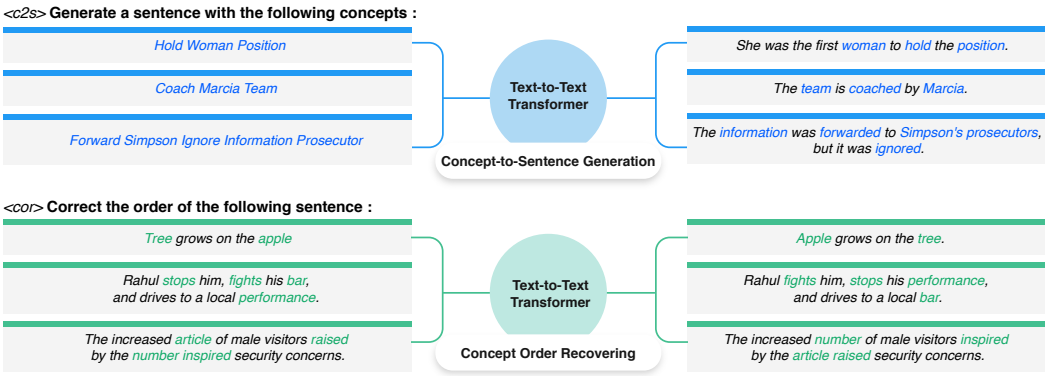

Figure 1: **Overview of Generative self-supervised pre-training objectives.** (1) **Concept-to-Sentence Generation** requires the text-to-text transformer to recover the original sentence given shuffled concept set (*Hold*, *Woman*, *Position* –> She was the first *woman* to *hold* the *position*). (2) **Concept Order Recovering** requires the text-to-text transformer to recover the original sentence given concept-shuffled sentence (*Tree* grows on the *apple* –> *Apple* grows on the *tree*).

We present two kinds of self-supervised pre-training tasks: **concept-to-sentence generation (C2S)** and **concept order recovering (COR)**. C2S trains the pre-trained model to compose ("write") sentences given a set of concepts, and expects the generated sentences to be fluent and plausible in terms of commonsense. COR aims to teach models to detect and revise a corrupted sentence with incorrect ordering of concepts. As illustrated in Figure 1, both tasks require a pre-trained model to recall relevant commonsense facts about the concepts and to understand the underlying commonsense relations between them. These pre-training objectives can together explicitly encourage the model to capture the relational concept-centric commonsense knowledge and perform compositional reasoning.

Specifically, we need a generative pre-training objective to encourage models to capture this **generative commonsense** reasoning ability, so that models can learn to generate sentences with commonsense knowledge for both C2S and COR. Also, to teach modes to distinguish truth sentences from less plausible ones, we need to teach models with **discriminative commonsense** through *contrastive self-training*. To unify both *generative* and *contrastive* objectives within a joint learning framework, we propose to use the generated sentences from the generative module as the distractors for learning the contrastive module. We share all the parameters between the generator (trained with the generative objective) and the discriminator (trained with the contrastive objective), then train multiple objectives with different prefixes. Compared with previous works (Peters et al., 2019; Li et al., 2019; Xiong et al., 2020) that utilize external knowledge bases like Wikidata or ConceptNet (Liu & Singh, 2004), our approach can directly improve the generative and discriminative commonsense reasoning ability of PTLMs at the same time without external knowledge bases.

To evaluate the effectiveness of our proposed method, we apply our method in an intermediate-task transfer learning setting (Pruksachatkun et al., 2020) based on the pre-trained T5-base model to train a *Concept-Aware Language Model* (**CALM**). While only continually pre-trained on a small dataset for a relatively fewer number of updates (compared to conventional pre-training), CALM consistently outperforms T5-base on four commonsense-related NLU datasets (i.e., COMMONSENSEQA, OPENBOOKQA, PIQA, and ANLI) and COMMONGEN, a commonsense-related NLG dataset. Our results and careful ablation studies demonstrate the potential of our method to serve as a "plug-and-play" method for any pre-trained text-to-text transformer before fine-tuning on commonsense-related tasks. To the best of our knowledge, our work is the first to investigate concept-centric self-supervised objectives that improve both generative and discriminative commonsense reasoning ability of a pre-trained language model.

2 SELF-SUPERVISED OBJECTIVES FOR CONCEPT-CENTRIC LEARNING

In this section, we first describe the proposed generative and contrastive objectives used for improving the commonsense reasoning ability of pre-trained text-to-text transformers. Then, we introduce the joint learning framework which unifies the proposed self-supervised objectives and learn a unified text-to-text transformer based on pre-trained models such as T5.

2.1 GENERATIVE OBJECTIVES

Similar to many other pre-training tasks such as masked language modeling, we aim to teach models to recover original sentences from corrupted inputs, which is often regarded as a *denoising* process. We propose two generative self-supervised pre-training objectives: concept-to-sentence generation (C2S) and concept order recovering (COR).

Concept Extraction. Given an input $\mathbf{x} = [x_1, x_2, \ldots, x_n]$, we first conduct part-of-speech tagging for the sentence and extract *Verb*, *Noun*, and *Proper Nouns* from the sentence to use as concepts[2]. Next, we form concept-sets $\mathcal{C} = [v_1, v_2, \ldots, v_p; n_1, n_2 \ldots, n_q]$ where v_i and n_i denotes the i-th verb or noun/proper noun concept (token) in $\mathbf{x}$. We denote $\mathcal{C}_v$ and $\mathcal{C}_n$ as the set of verb and noun/proper noun concepts respectively in $\mathcal{C}$. (i.e. $\mathcal{C}_v = [v_1, v_2, \ldots, v_p]$ and $\mathcal{C}_n = [n_1, n_2, \ldots, n_q]$.)

Concept-to-Sentence Generation (C2S). The concept-to-sentence generation (C2S) objective requires the text-to-text transformer to recover the original sentence given only a few unordered keywords of the sentence. Specifically, given a sentence, we shuffle the extracted concept-set $\mathcal{C}$ to create the perturbed source sequence and train the model to generate the original sentence with a prefix (denoted as <c2s>) as described in Fig. 1. Formally, the C2S objective can be formulated as:

$$L_{c2s} = \mathbb{E}\Big(\sum_{i=1}^{n} -\log p(x_i|\texttt{<c2s>}; \textsc{permute}(\mathcal{C}); x_{1:i-1})\Big) \tag{1}$$

where the PERMUTE() function randomly shuffle the concepts in the concept-set. This objective requires the model to construct an acceptable commonsense sentence by adhering to and reasoning over the commonsense relations between the given concepts. Therefore, relational commonsense knowledge is implicitly injected into the parameters of the model. The C2S objective is motivated by the task proposed in Lin et al. (2019). Compared to their work, the concept-set used in C2S covers more concepts such as named entities, while the original task only includes the concepts appearing in ConceptNet. We apply the task in a general domain and as a pre-training objective, instead of merely serving as an evaluation task.

Concept Order Recovering (COR). As for the concept order recovering (COR) objective, we shuffle the order of concept in a sentence and train the model to recover the original sentence. As illustrated in Figure 1, given an input sentence "tree grows on the apple,", the models would shuffle the concepts including "tree", "grow", and "apple" to recover the original sentence "apple grows on the tree." The noise introduced by concept shuffling is different from that by traditional self-supervised objectives like mask language modeling and mask span prediction because the corrupted source sentences are in general complete (i.e., no tokens or spans are masked) and grammatically correct, while not acceptable in terms of commonsense because the order and relation between concepts are shuffled. By training the model to detect and correct the disorder of concepts in a sentence, the model is expected to acquire some relational commonsense knowledge like "apple generally grows on a tree" instead of "tree grows on an apple."

Formally, the COR objective can be formulated as:

$$L_{cor} = \mathbb{E}\Big(\sum_{i=1}^{n} -\log p(x_i|\texttt{<cor>}; \textsc{concept-permute}(\mathbf{x}, \mathcal{C}); x_{1:i-1})\Big), \tag{2}$$

where <cor> is the prefix for the COR objective illustrated in Figure 1. The function CONCEPT-PERMUTE() permutes the order between concepts in the same category (i.e. noun or verb) in the sentence, which can be formally defined as:

$$\textsc{concept-permute}(\mathbf{x}, \mathcal{C}) = [x'_1, x'_2, \ldots, x'_n] \text{ where } x'_i = \begin{cases} x_i & x_i \notin \mathcal{C} \\ \textsc{permute}(\mathcal{C}_v)[j] & x_i = v_j \\ \textsc{permute}(\mathcal{C}_n)[j] & x_i = n_j \end{cases} \tag{3}$$

The key difference between our proposed objectives and others is that our objectives requires the model to capture the relational commonsense knowledge and perform compositional commonsense reasoning in order to successfully reconstruct the original sentence; while traditional pre-training objectives mainly focuses on general token-level patterns.

[2]We split the concepts with multiple tokens (under Spacy tokenization) into single token to ensure the concepts discussed afterwards all contain a single token.

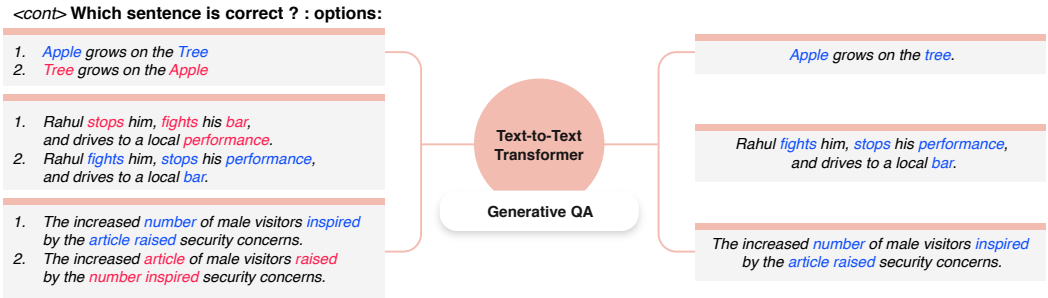

Figure 2: **Overview of Contrastive self-supervised pre-training objectives.** Generative QA style contrastive objective requires model to distinguish truth sentences from less plausible ones.

2.2 Contrastive Objective

The contrastive objective encourages the pre-trained model to distinguish the real sentence from a distractor sentence: a sentence that is similar to the real sentence, generally grammatically correct, but may not follow commonsense. We expect it to improve the pre-trained model's discriminative commonsense reasoning ability so that the model's performance on commonsense-reasoning-discriminative tasks, like CommonsenseQA, can be improved. We formulate the contrastive objective as a **Generative QA** task: we take the concatenation of a prefix <cont> (question / context), the real sentence x (answer), and the distractor x' (distractor) as the input and train the model to output the real sentence x. Formally, we have the loss function of the contrastive objective defined as:

$$L_{cont} = \mathbb{E}\big(-\log p(x|\texttt{<cont>}; \textsc{permute}(x; x'))\big), \tag{4}$$

where the prefix <cont> is described in Figure 2. The distractor x' is either constructed by concept shuffling as described previously (i.e. $x' = \textsc{concept-permute}(\mathbf{x}, \mathcal{C})$) when used independently, or generated by a generator trained with the aforementioned generative objectives when used in the joint training framework, which will be described in the next section.

3 Joint Training with Generative and Contrastive Objectives

The aforementioned generative and contrastive self-supervised objectives can be applied independently or simply combined in a multi-task learning fashion. We argue that these two objectives can mutually reinforce each other: the generated sentences from the generative objective can help the contrastive module learn to distinguish commonsense sentences from less plausible ones. Therefore, we propose a joint training framework to unify generative objectives and contrastive objectives by using the generator to produce *distractors* for learning towards contrastive objective.

Specifically, we have a **generator** G_θ (trained with the generative objectives) and a **discriminator** D_ϕ (trained with the contrastive objective). Given an input sentence $\mathbf{x}$, we first use the method for either C2S or COR to produce the corrupted source sequence $\mathbf{x}'$. Then, we use the generator G_θ trained with the corresponding objective to generate the recovered sentence $\mathbf{x}'' = G_\theta(x')$. We then take $\mathbf{x}''$ as the distractor to train the discriminator D_ϕ with the contrastive objective. The loss function of the proposed joint training framework consists of two parts: the first part is the loss of generative objectives, which is identical to the loss described in Eq.(1) and Eq.(2) and is used to update the generator G_θ. The second part is the loss of the contrastive objective as described in Eq.(4), which can be formulated as:

Algorithm 1: Pre-training Concept-Aware Language Model (CALM).

Input: Text-to-Text Transformer T_θ, Text corpus X=[$x_1, x_2, \ldots, x_n$].
repeat
  for *each* $x_i \in X$ **do**
    Extract the concept-set $\mathcal{C}_i$;
    Construct the distractor sentence $x' = \textsc{concept-permute}(\mathbf{x}_i, \mathcal{C}_i)$;
    Update T_θ with Eq.(1, 2, 4);
until *maximum iterations reached*;
repeat
  for *each* $x_i \in X$ **do**
    Update T_θ with Eq.(7)
until *maximum iterations reached*;

$$L_{cont_joint_c2s} = \mathbb{E}\big(-\log \mathrm{D}_\phi(y|\texttt{<cont>}; x; \mathrm{G}_\theta(\texttt{<c2s>}; \textsc{permute}(\mathcal{C}))\big) \tag{5}$$

$$L_{cont_joint_cor} = \mathbb{E}\big(-\log \mathrm{D}_\phi(y|\texttt{<cont>}; x; \mathrm{G}_\theta(\texttt{<cor>}; \textsc{concept-permute}(\mathbf{x}, \mathcal{C}))\big) \tag{6}$$

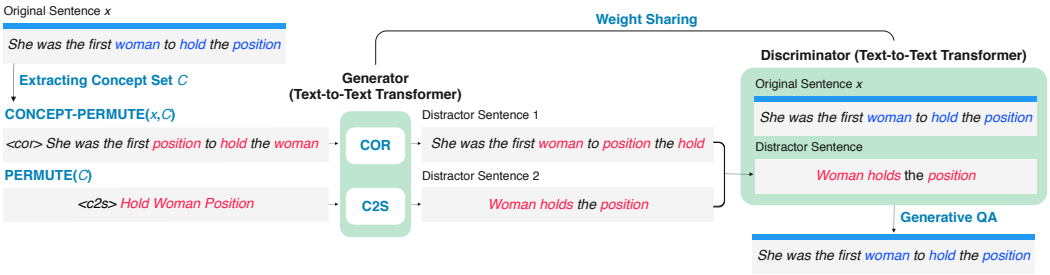

Figure 3: **Proposed Joint Training Framework.** Given an input sentence x ("She was the first *woman* to *hold* the *position*."), we extract concept-set $\mathcal{C}$ (*woman*, *hold*, *position*). Given x and $\mathcal{C}$, we produce corrupted source sequence x' either for C2S and COR. The generator trained with the corresponding objective recovers sentences as distractors x'' to the discriminator. The discriminator is trained to distinguish truth sentences from randomly selected distractor among two objectives. Parameters between the generator and discriminator are shared.

where $L_{cont_joint_c2s}$ and $L_{cont_joint_cor}$ is the contrastive loss with the distractor generated with either the C2S or the COR objective. We then have the overall objective for the joint training framework defined as :

$$L_{joint} = (L_{c2s} + L_{cor}) + \beta(L_{cont_joint_c2s} + L_{cont_joint_cor}). \tag{7}$$

L_{c2s} and L_{cor} are defined in Eq.(1) and Eq.(2) respectively and β is a hyperparameter controlling the relative weight between the generative and contrastive objectives. Note that since we would like to inject both generative and discriminative commonsense reasoning ability into the parameters of a single text-to-text transformer, we share the parameters between the generator G_θ and the discriminator D_ϕ.

Finally, we describe the overall procedure to apply the proposed self-supervised objectives and the joint training framework on a pre-trained text-to-text transformer. We apply a two-stage training strategy. During the first stage, we apply our proposed generative and contrastive objectives individually on the model in a multi-task learning fashion with different prefixes. This provides a good starting point for the second stage where the joint training framework is applied. We summarize the workflow of our method in Algorithm 1.

4 Experiments

In this section, motivated by the observation of Pruksachatkun et al. (2020) that tasks requiring commonsense reasoning ability generally serve as good intermediate task, we test our method in the intermediate task transfer setting. Specifically, we initialize our model with T5-base, a pre-trained text-to-text transformer model, and training the model with our proposed method as intermediate task before fine-tuning and target downstream tasks. Another reason for adopting this setting is because we expect our method to serve as a "plug-and-play" method that can be applied to any pre-trained text-to-text transformer by simply continually training for a few steps.

Details for Pre-training and Fine-tuning CALM is continually pre-trained with our proposed self-supervised objectives as intermediate tasks based on the pre-trained T5-base model following the setting in Pruksachatkun et al. (2020). We randomly sample 500K sentences from the English Wikipedia corpus, which is used for pre-training BERT and its variants, as the source dataset for our proposed self-supervised objectives which serve as intermediate tasks. We then fine-tune the CALM on each downstream task individually and report the average performance of three runs with different random seeds for fine-tuning on each dataset since the performance is sensitive to different random seeds. Training details and hyperparameter settings are presented in Appendix A.1 and A.2.

Datasets We consider five commonsense benchmark datasets as target tasks. We categorize these datasets into discriminative and generative tasks. Discriminative tasks are classification tasks while generative tasks are text generation tasks. We consider four datasets for discriminative task: **CommonsenseQA** (Talmor et al., 2018), **OpenbookQA** (Mihaylov et al., 2018), **PIQA** (Bisk et al., 2020), **aNLI** (Bhagavatula et al., 2019) and one dataset for generative task: **CommonGEN** (Lin et al., 2019). Details on datasets are discussed in Appendix A.3.

Methods	CSQA	OBQA	PIQA	aNLI	CommonGEN			
	Accuracy				BLEU-4	METEOR	CIDEr	SPICE
T5-base	61.88(±0.08)	58.20(±1.0)	68.14(±0.73)	61.10(±0.38)	24.90	31.20	12.99	32.40
T5-base w/ additional epochs	61.92(±0.45)	58.10(±0.9)	68.19(±0.77)	61.15(±0.52)	25.10	31.00	13.12	32.40
T5-base + SSM	62.08(±0.41)	58.30(±0.8)	68.27(±0.71)	61.25(±0.51)	25.20	31.20	13.28	32.40
CALM (Generative-Only)	62.28(±0.36)	58.90(±0.4)	68.91(±0.88)	60.95(±0.46)	25.80	31.20	13.81	32.60
CALM (Contrastive-Only)	62.73(±0.41)	59.30(±0.3)	70.67(±0.98)	61.35(±0.06)	25.50	31.20	13.58	32.60
CALM (Mix-only)	63.02(±0.47)	60.40(±0.4)	70.07(±0.98)	62.79(±0.55)	26.00	31.20	13.82	32.80
CALM (w/o Mix warmup)	62.18(±0.48)	59.00(±0.5)	69.21(±0.57)	61.25(±0.55)	25.80	31.20	13.77	32.60
CALM	**63.32(±0.35)**	**60.90(±0.4)**	**71.01(±0.61)**	**63.20(±0.52)**	**26.40**	**31.40**	**13.88**	**33.00**

Table 1: **Experimental results on commonsense reasoning datasets.** The first group of models are baselines. The models in the middle group are trained with the proposed objectives independently and the last group of models are trained by joint training. Best models are bold and second best ones are underlined within each metric.

Compared Methods We compare our model with following models continually trained with different intermediate tasks based on the pre-trained T5-base model: (1) **T5-base** is the pre-trained T5-base model without continually training on any intermediate task. (2) **T5-base w/ additional epochs** is continually pre-trained using the original pre-training objective of T5 with additional training steps. The total number of additional training steps is equal to that of our final model. (3) **T5-base + SSM** is continual pre-trained with a variant of the *salient span masking* objective (Guu et al., 2020; Roberts et al., 2020) objective that masks text spans of concepts extracted with POS tagging instead of named entities extracted by a pre-trained NER model, which makes it more focused on concepts. (4) **CALM(Generative-Only)** is continually pre-trained with the proposed generative objectives including concept-to-sentence generation(C2S) and concept order recovering(COR) as intermediate tasks. (5) **CALM(Contrastive-Only)** is continually pre-trained with the proposed contrastive objective as described in section 2.2 using the distractor generated by concept shuffling. (6) **CALM(Mix-only)** is continually pre-trained with both the generative objectives and the contrastive objective, combined with a multi-task learning fashion[3] with identical weights for each objective as the intermediate task. (7) **CALM (w/o Mix warmup)** is continually pre-trained with the joint training objective described in Eq (7) directly from the pre-trained T5-base model. (8) **CALM** is our main model trained as described in Algorithm 1. The difference between CALM and CALM (Joint) is that the former is initialized by the CALM(Mix).

Evaluation Metrics For discriminative tasks, we choose accuracy as our metric following other conventional question answering tasks. For generative tasks, we report automated metrics including BLEU (Papineni et al., 2002), METEOR (Banerjee & Lavie, 2005), CIDEr (Vedantam et al., 2015), and SPICE (Anderson et al., 2016) following the leaderboard of COMMONGEN (Lin et al., 2019). Results for COMMONGEN are on the test set and others are on the official development set. We tune the hyperparameters based on the models' performance on a in-house split dev set.

4.1 EXPERIMENTAL RESULTS

The result is presented in Table 1. First, we can see that our CALM model consistently outperforms the backbone T5-base model on all five datasets by a margin range from 1.5 to 2.9 accuracy on discriminative tasks and 1.5/0.6 BLEU/SPICE score on CommonGEN. This is an impressive result since we are only performing intermediate training on a relatively small dataset for only around 20k updates. It demonstrates the potential of our method for serving as a "plug-and-play" method for packing more commonsense knowledge into a pre-trained text-to-text transformer. We also find that CALM performs comparably with several large-size PTLMs like BART, T5-large, and GPT-2 on the COMMONGEN dataset and is only slightly worse than KG-BART (Liu et al., 2020), the current

Methods	Params	CommonGEN			
		BLEU-4	METEOR	CIDEr	SPICE
GPT-2 (Radford et al., 2019)	774M	21.10	26.20	12.15	25.90
UniLM (Dong et al., 2019)	340M	27.70	29.70	14.85	30.20
BART (Lewis et al., 2020)	406M	26.30	30.90	13.92	30.60
T5-Base (Raffel et al., 2019)	220M	16.40	23.00	9.16	22.00
T5-Large (Raffel et al., 2019)	770M	28.60	30.10	14.96	31.60
KG-BART (Liu et al., 2020)	406M	**30.90**	**32.40**	**16.83**	32.70
Our T5-Base	220M	24.90	31.20	12.99	32.40
CALM	220M	26.40	31.40	13.88	**33.00**

Table 2: Comparison between PTLMs on CommonGEN. Above baselines are reported number in the leaderboard. Best models are bolded.

[3]We tried combining them by training with the generative and contrastive objectives sequentially but the preliminary results are much worse than directly combining them with multi-task learning.

	CSQA	PIQA	COMMONGEN			
Method	Accuracy		BLEU-4	METEOR	CIDEr	SPICE
T5 - Text Infilling	61.92	68.19	25.10	31.00	13.13	32.40
CALM - COR	**62.36**	**68.77**	25.70	31.20	13.65	32.60
CALM - C2S	62.24	68.75	**25.90**	**31.40**	**13.94**	**32.80**

(a) Generative objectives

	CSQA	PIQA	COMMONGEN			
Metric	Accuracy		BLEU-4	METEOR	CIDEr	SPICE
Multi-choice QA	62.21	68.82	25.00	31.20	13.28	32.60
True/False	62.24	67.81	25.10	31.20	13.41	32.60
Generative QA	**62.73**	**70.67**	**25.50**	31.20	**13.58**	32.60

(b) Contrastive objectives

Table 3: **Analysis on Contrastive and Generative objectives.**

state-of-the-art on COMMONGEN, which is a contemporary work that exploits external knowledge bases as additional information, and is based on a larger backbone(i.e., BART (Lewis et al., 2019)).

In addition, we can observe that both the proposed generative and contrastive objective outperforms the backbone T5-base model, as well as its variants that continually pre-trained with the original masked span prediction objective and the concept-specific salient span masking scheme, when applied independently. The combination of the generative and contrastive objectives (i.e., CALM(Mix-only)) yields further improvement upon the model trained independently with either generative or contrastive objectives. Also, we find that the CALM model consistently outperforms CALM(Mix), demonstrating the effectiveness of the proposed joint training framework. Applying joint training directly on top of a pre-trained model (i.e., CALM(w/o Mix warmup)) does not work very well, demonstrating the necessity of applying mixed training to initialize the model before starting joint training.

4.2 PERFORMANCE ANALYSIS

Analysis on Generative objective To investigate the contribution of each generative objective, we conduct an ablation study by continually pre-training three models from the same T5-base model with C2S, COR, and text infilling, which is the original objective for pre-training T5, as the objective for the intermediate task. We continually pre-train these models for the same number of steps and then evaluate their performance by fine-tuning on different target tasks. The result is shown in Table 3. We can see that both C2S and COR works better than the original masked span infilling objective on itself. This confirms the effectiveness of our proposed generative objectives on improving the commonsense reasoning ability of pre-trained text-to-text transformers.

Task Formulation of the Contrastive objectives For contrastive objectives, we test three different task formats: Multi-choice QA, Generative QA, and True/False. Multi-choice QA and Generative QA takes the concatenation of the real sentence and the distractor. Then, Multi-choice QA output the index of the real sentence following other conventional Multi-choice QA tasks, and Generative QA output the real sentence respectively. True/False takes either the real sentence or the distractor and train the model to perform a binary classification problem of whether the input sentence makes sense. The result is shown in Table 3. We could find that the format of Generative QA performs the best. We suspect this is because the Generative QA format is closer to the format used during the original pre-training stage of the T5 model and the format used for fine-tuning.

Performance with fewer training examples To investigate the effectiveness of our objective in the low-resource setting, we explore the performance of our model and baselines fine-tuning with different fractions of the training data. From Figure 4, we can see that the performance improvement yielded by our models upon the T5-base model is more significant in the low-resource regime. This shows that CALM may already pack some commonsense knowledge in its parameters so that it does not require much data for fine-tuning before obtaining a good performance. In contrast, the original T5 model requires much data for fine-tuning, which suggests it may fail to encode much commonsense knowledge and must fit the correlation patterns in the downstream datasets to get a good performance.

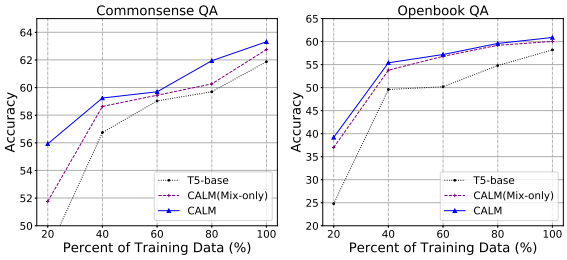

Figure 4: Performance of compared models fine-tuned with different fraction of the datasets.

Comparison of Generated Data Table 4 shows the comparison of generated examples for the COMMONGEN test set between T5-base and CALM. We can see that the sentences generated by CALM are generally more acceptable in terms of commonsense plausibility while T5-base sometimes generates sentences that do not make sense.

Concept-set	T5-base	CALM
Grass, Dog, Ball, Chase	a dog is chased by a ball on the grass.	dog chasing a ball in the grass.
Net, Cast, Boat, Water	fishing boat casts a net in the water.	fisherman casts a net into the water from a fishing boat.
Hole, Tree, Plant, Dig	a man digs a hole in a tree to plant a new tree . he digs the	man digging a hole to plant a tree.
Ingredient, Add, Pan, Fry	a pan filled with ingredients adds a touch of spice to the fry .	add the ingredients to a pan and fry.
Water, Hold, Hand, Walk	A man holding a hand and walking in the water. A man is holding water.	man holding a bottle of water in his hand as he walks down the street.
Place, Use, Metal tool	A man uses a metal tool to make a piece of metal.	woman uses a metal tool to make a piece of jewelry.
Hair, Wax, Apply, Remove	remove the wax from your hair and apply it to your hair .	woman applies wax to her hair and then removes it with a comb.
Sidewalk, Dog, Walk, Leash	A dog walking on a leash on the sidewalk.	dog walking on a sidewalk with a leash.

Table 4: **Comparison of generated sentences with same concept-set**

5 Related Work

Self-Supervised Language Representation Pre-Training. Motivated by the fact that words can have different meanings in different contexts, contextual language representation methods (McCann et al., 2017; Peters et al., 2018) have been developed and shown superior performance on downstream tasks compared with static word embeddings Mikolov et al. (2013); Pennington et al. (2014). More recently, large scale language models based on transformer architecture (Vaswani et al., 2017) pre-trained with either mask language modeling objective (Devlin et al., 2018; Liu et al., 2019; Lan et al., 2019) or mask span infilling objective (Lewis et al., 2019; Raffel et al., 2019) have been explored further advanced the state-of-the-art on multiple NLU and NLG tasks. Our method is based on these techniques and we focus on improving the commonsense reasoning ability of pre-trained text-to-text transformers. More recently, Clark et al. (2020) propose a new self-supervised pre-training objective called Replaced Token Detection (RTD). RTD uses a mask language model like BERT to fill in the mask and train a discriminator to predict whether a token is generated or real. This pre-training paradigm is related to our proposed joint training framework. Some major differences include that (1) Our method employs sentence-level distractors that are in general grammatically correct but not in line with commonsense, thus require the model to perform relational commonsense reasoning while RTD is a token-level discrimination task and can often be solved with syntactic and shallow semantic knowledge (Rosset et al., 2020); (2) Our method unifies generative and contrastive objectives with one model, which can be applied to both NLU and NLG downstream tasks; and (3) The discriminator in our framework is "contrastive", takes both the real sentence and the distractor as input simultaneously.

Knowledge-augmented PTLMs. As standard pre-trained language models usually do not explicitly model knowledge, a number of works have examined the problem of incorporating world knowledge with the PTLMs. Recent work Zhang et al. (2019); Peters et al. (2019); Wang et al. (2020); Liu et al. (2020) utilizes an external knowledge base to incorporate entity knowledge with PTLMs; however, these approaches require specialized resources like knowledge bases, which limits the domain they can be applied to. Xiong et al. (2020) proposes WikiLM that encodes world knowledge into the parameters of a BERT(Devlin et al., 2018)-like pre-trained model with a novel entity replacement detection objective that incorporates Wikipedia to form distractors. Their approach differs from ours because it requires an external knowledge base (i.e., Wikipedia) which limits the domain it can be applied, is limited to discriminative pre-training objectives and downstream tasks, and focuses on world knowledge instead of relational commonsense knowledge. More recently, (Rosset et al., 2020) propose KALM, an entity-aware language model with more world knowledge packed into its parameters. Their method is restricted to the training of language models instead of masked language models or text-to-text transformers which can be used for more downstream tasks. Also, all the aforementioned work mainly focuses on world knowledge of named entities. In contrast, our work mainly focuses on commonsense knowledge about quotidian concepts.

6 Conclusion

We propose novel self-supervised strategies that encourage the model to focus on concept-centric information that is related to commonsense understanding and reasoning instead of simple word co-ocurrence patterns so that the commonsense learning capability of pre-trained text-to-text transformers can be improved. Despite merely continually pre-trained on a small dataset with only around 20k steps, our CALM model consistently outperforms the T5-base model on all commonsense-related datasets, and even yields better performance compared with some larger size PTLMs on the CommonGEN dataset. The performance gain is larger when we use fewer examples for fine-tuning on different downstream tasks, indicating that CALM effectively encodes more commonsense knowledge and rely less on fitting superficial patterns of datasets compared to traditional pre-trained language models. Our work suggests that text-to-text models can be pre-trained with better parameter and sample efficiency by carefully designed self-supervised objectives that focus more on the ability (e.g., commonsense reasoning ability) required by target tasks.

REFERENCES

Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *ECCV*, 2016.

Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pp. 65–72, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics. URL `https://www.aclweb.org/anthology/W05-0909`.

Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen-tau Yih, and Yejin Choi. Abductive commonsense reasoning. *arXiv preprint arXiv:1908.05739*, 2019.

Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about physical commonsense in natural language. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.

Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020.

Ernest Davis and Gary Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM*, 58:92–103, 08 2015. doi: 10.1145/2701413.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. Unified language model pre-training for natural language understanding and generation. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 13063–13075. 2019.

WA Falcon. Pytorch lightning. *GitHub. Note: https://github.com/PyTorchLightning/pytorch-lightning*, 3, 2019.

Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. Realm: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*, 2020.

D. Khashabi, S. Min, T. Khot, A. Sabhwaral, O. Tafjord, P. Clark, and H. Hajishirzi. Unifiedqa: Crossing format boundaries with a single qa system. *EMNLP - findings*, 2020.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 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*, Online, July 2020. Association for Computational Linguistics.

Shiyang Li, Jianshu Chen, and Dian Yu. Teaching pretrained models with commonsense reasoning: A preliminary kb-based approach. *arXiv preprint arXiv:1909.09743*, 2019.

Bill Yuchen Lin, Wangchunshu Zhou Ming Shen, , Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. Commongen: A constrained text generation challenge for generative commonsense reasoning. *CoRR*, abs/1911.03705, 2019.

Hugo Liu and Push Singh. Conceptnet—a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226, 2004.

Ye Liu, Yao Wan, Lifang He, Hao Peng, and Philip S. Yu. Kg-bart: Knowledge graph-augmented bart for generative commonsense reasoning, 2020.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. Learned in translation: Contextualized word vectors. In *Advances in Neural Information Processing Systems*, pp. 6294–6305, 2017.

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL `https://www.aclweb.org/anthology/P02-1040`.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.

Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 2018.

Matthew E Peters, Mark Neumann, Robert L Logan IV, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. Knowledge enhanced contextual word representations. *arXiv preprint arXiv:1909.04164*, 2019.

Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R Bowman. Intermediate-task transfer learning with pretrained models for natural language understanding: When and why does it work? *arXiv preprint arXiv:2005.00628*, 2020.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.

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. *arXiv preprint arXiv:1910.10683*, 2019.

Adam Roberts, Colin Raffel, and Noam Shazeer. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*, 2020.

Corby Rosset, Chenyan Xiong, Minh Phan, Xia Song, Paul Bennett, and Saurabh Tiwary. Knowledge-aware language model pretraining. *arXiv preprint arXiv:2007.00655*, 2020.

Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4566–4575, 2015.

Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Cuihong Cao, Daxin Jiang, Ming Zhou, et al. K-adapter: Infusing knowledge into pre-trained models with adapters. *arXiv preprint arXiv:2002.01808*, 2020.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771, 2019.

Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=BJlzm64tDH`.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. Ernie: Enhanced language representation with informative entities. *arXiv preprint arXiv:1905.07129*, 2019.

A APPENDIX

A.1 PRE-TRAINING DETAILS

The following details apply to both base architecture and joint-training architecture. We implement our pre-train models using Pytorch-lightning (Falcon, 2019) and Hugginface's Pytorch Transformers (Wolf et al., 2019). For pre-training phase, we use the Adam optimizer with maximum sequence length 256, train batch size 8, gradient accumulation 8, warmup steps 10000, weight decay 0.01 and adam epsilon 1e-6. We train the models with 8 V100 GPUs and FP32 precision for 17 hours. The model is pre-trained for at most 3 epochs to prevent overfitting. We searched for the best learning rate for our model out of [1e-4, 2e-5, 2e-6, 5e-7].

A.2 FINE-TUNING DETAILS

For fine-tuning, we use 4 V100 GPUs and use FP32. For all discriminative tasks, we use the Adam optimizer with maximum sequence length 256, batch size 4 and gradient accumulation 16. For generative task, we use the Adam optimizer with maximum source length 32, maximum target length 32, batch size 8, gradient accumulation 16. For all tasks, we use warmup fraction 0.01. Learning rates and train epochs are listed in Table 5.

Hyperparameter	CommonsenseQA	OpenbookQA	PIQA	aNLI	CommonGEN
Learning rate	[1e-4, 2e-4, 3e-4]	[5e-5, 1e-4, 2e-4, 3e-4]	[1e-4, 2e-4, 3e-4]	[2e-5, 3e-5]	[2e-5]
Train Epochs	20	20	20	10	20

Table 5: **Fine-tuning hyperparameters.**

A.3 DATASET PROPERTIES

- **CommonsenseQA** (Talmor et al., 2018) is a multiple-choice question answering task, which picks the most appropriate answer on general commonsense questions.
- **OpenbookQA** (Mihaylov et al., 2018) is a multiple-choice question answering task, which is modeled after open book exams on elementary-level core science questions. The task requires open book fact and additional commonsense which is not contained in the book. To test the commonsense reasoning ability, we do not use open book fact.
- **PIQA** (Bisk et al., 2020) is multiple-choice question answering task, which chooses the most appropriate solution for physical commonsense questions.
- **aNLI** (Bhagavatula et al., 2019) is a binary-classification task, which picks the most plausible explanatory hypothesis given two observations from narrative contexts.
- **CommonGEN** (Lin et al., 2019) is a constrained text generation task, which generates a coherent sentence describing an everyday scenario using common concepts.

Dataset	Train	Development	Test	Source Example	Target Example
CommonsenseQA	9,741	1,221	1,140	context: *What home entertainment equipment requires cable?* options: 1: *radio shack* 2: *substation* 3: *cabinet* 4: *television* 5: *desk*	4
OpenbookQA	4,957	500	500	context: *You can make a telescope with* options: 1: *straw* 2: *glass* 3: *candle* 4: *mailing tube*	2
PIQA	16,113	1,838	3,084	context: *When boiling butter, when it's ready, you can* options: 1: *Pour it onto a plate* 2: *Pour it into a jar*	2
aNLI	169,654	1,532	3,040	context: *It was my birthday. When I got home the party was set up for my brother.* options: 1: *I was so excited.* 2: *I was so mad.*	2
CommonGEN	67,389	4,018	6,042	generate a sentence with these concepts: *Apple Grow Tree*	*Apple grows* on the *tree*

Table 6: **Properties of Commonsense benchmark datasets.**