

Generative Data Augmentation for Commonsense Reasoning

Yiben Yang[♣] Chaitanya Malaviya[†]^{II} Jared Fernandez[♣][♥] Swabha Swayamdipta[†]
Ronan Le Bras[†] Ji-Ping Wang[♣] Chandra Bhagavatula[†] Yejin Choi[†][◇] Doug Downey[♣][†]

[♣]Northwestern University, Evanston, IL, USA

[†]Allen Institute for Artificial Intelligence, Seattle, WA, USA

^{II} University of Pennsylvania, Philadelphia, PA, USA

[♥] Carnegie Mellon University, Pittsburgh, PA, USA

[◇]University of Washington, Seattle, WA, USA

{yiben.yang@, jared.fern@u., jzwang@}northwestern.edu

{chaitanyam, swabhas, ronanlb, chandrab, yejin, dougd}@allenai.org

Abstract

Recent advances in commonsense reasoning depend on large-scale human-annotated training sets to achieve peak performance. However, manual curation of training sets is expensive and has been shown to introduce annotation artifacts that neural models can readily exploit and overfit to. We propose a novel generative data augmentation technique, **G-DAUG^c**, that aims to achieve more accurate and robust learning in a low-resource setting. Our approach generates synthetic examples using pretrained language models, and selects the most informative and diverse set of examples for data augmentation. On experiments with multiple commonsense reasoning benchmarks, G-DAUG^c consistently outperforms existing data augmentation methods based on back-translation, establishing a new state-of-the-art on WINOGRANDE, CODAH, and COMMONSENSEQA, and also enhances out-of-distribution generalization, proving to be more robust against adversaries or perturbations. Our analysis demonstrates that G-DAUG^c produces a diverse set of fluent training examples, and that its selection and training approaches are important for performance.

1 Introduction

While recent advances in large-scale neural language models (Devlin et al., 2019; Liu et al., 2019; Radford et al., 2019; Raffel et al., 2019) have led to strong performance on several commonsense reasoning benchmarks (Talmor et al., 2019; Lv et al., 2020; Sakaguchi et al., 2020), their accuracy by and large depends on the availability of large-scale human-authored training data. However, crowdsourcing examples at scale for each new task and domain can be prohibitively expensive. Moreover, human-authored data has been shown to exhibit annotation artifacts (Gururangan et al., 2018; Agrawal

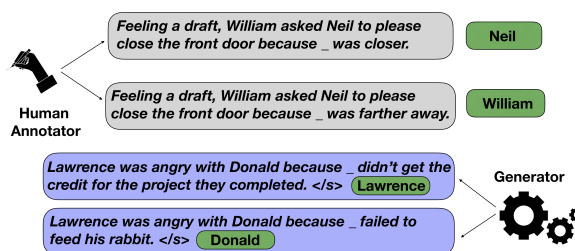


Figure 1: Example of a selected high-quality generated example compared to a human-authored example from the WINOGRANDE dataset. Composing commonsense questions can require creativity.

et al., 2018; Schwartz et al., 2017), leading to models with considerably weaker performance on out-of-distribution samples (Jia and Liang, 2017; Belinkov and Bisk, 2017; Iyyer et al., 2018).

A candidate solution that has shown promise in other tasks, such as reading comprehension, is to augment a human-authored training set with a large set of synthetically-generated examples (Zhou et al., 2017; Du et al., 2017; Zhao et al., 2018a). But, generating synthetic examples for commonsense reasoning poses a unique challenge. In reading comprehension, for instance, the goal of data augmentation is to generate questions that are directly answerable by a given reference passage. In contrast, answering commonsense questions relies on commonsense notions that are seldom stated explicitly (Gordon and Van Durme, 2013; Forbes and Choi, 2017), and authoring such questions can require creativity (see Figure 1). Based on promising evidence from previous work (Yang et al., 2018; Trinh and Le, 2018; Bosselut et al., 2019; Davison et al., 2019), we hypothesize that pretrained language models, such as GPT-2 (Radford et al., 2019), capture some common sense expressed implicitly in their pretraining corpus. Could questions generated by such models serve as helpful

training data? In this work, we explore this question through **Generative Data Augmentation** for commonsense reasoning (**G-DAUG^c**; §2): a novel framework for augmenting training data with *diverse* and *informative* synthetic training examples to improve both in-distribution performance and out-of-distribution generalization of commonsense reasoning models.¹

Although a generative model allows us to produce large pools of synthetic training examples, the generated examples may be noisy or redundant. To ensure that we use the most informative examples for augmentation, we introduce data selection methods based on influence functions (Koh and Liang, 2017) and a heuristic to maximize the diversity of the generated data pool. Finally, we propose an effective two-stage training scheme for augmentation with synthetic data. In experiments across multiple commonsense benchmarks, we show that G-DAUG^c can mitigate the expense and brittleness resulting from large training sets for commonsense reasoning tasks.

To summarize, our contributions include:

1. G-DAUG^c, a generative data augmentation framework for commonsense reasoning (§2),
2. novel selection methods that identify *informative* and *diverse* synthetic training examples from the generated pool (§3),
3. experiments showing that G-DAUG^c improves in-distribution performance, achieving a 1–4% average absolute gain across four commonsense reasoning data sets and state-of-the-art results on the WINOGRANDE (Sakaguchi et al., 2020), COMMONSENSEQA (Talmor et al., 2019), and CODAH (Chen et al., 2019) benchmarks, and also improves model robustness in terms of resistance to adversarial attacks (Jin et al., 2020) and accuracy on perturbed evaluation sets (§4), and
4. a comprehensive analysis of the factors that influence G-DAUG^c’s performance (§5).

2 G-DAUG^c

We now describe our framework for **Generative Data Augmentation for Commonsense Reasoning (G-DAUG^c)**. Figure 2 shows an overview of the approach. We describe G-DAUG^c’s data generation procedure (steps 1 and 2 in the figure) in this section, and cover the data selection and training

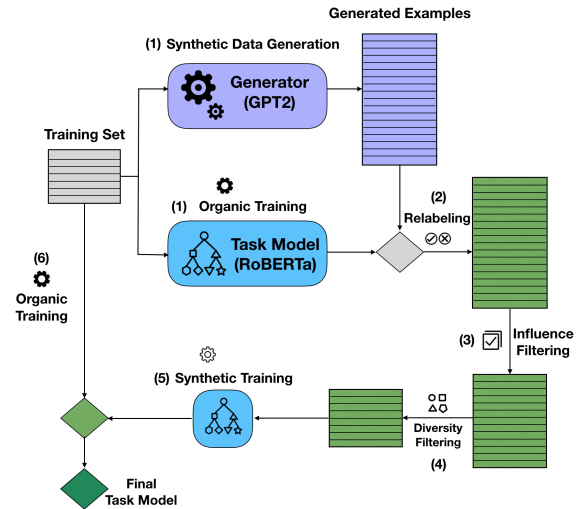


Figure 2: Illustration of the G-DAUG^c process: (1) generate synthetic data and train a task model, (2) relabel the generated data using the task model, (3) filter the generated data based on estimated influence scores, (4) further select a subset based on a diversity-maximizing heuristic, (5) train a new task model using the filtered generations (*synthetic training*), and (6) further train this model using the original training data (*organic training*).

components (steps 3-5) in §3.

2.1 Synthetic Training Data Generation

We will use multiple choice question answering as a running example to describe synthetic data generation. Formally, consider a dataset of N questions $\mathcal{D} = \{(\mathbf{Q}^i, \mathcal{C}^i, y^i) : i = 1, 2, \dots, N\}$, where \mathbf{Q}^i is a sequence of words denoting the i^{th} question, $\mathcal{C}^i = \{\mathbf{C}_j^i : j = 1, 2, \dots, K\}$ is the corresponding choice set with K choices which are word sequences as well, and a ground truth label $y^i \in \{1, 2, \dots, K\}$. We denote the **answer** as $\mathbf{C}_{y^i}^i$ and the **distractors** as $\mathbf{C}_{j \neq y^i}^i$ ’s.

Our text generators are pretrained generative language models, finetuned to maximize the log-likelihood of a sequence of text \mathbf{W} , $\mathcal{L}_W(\theta) = \sum_{t=1}^T \log P(w_t | \mathbf{W}_{1:t-1}; \theta)$, where $\mathbf{W}_{1:t-1}$ denotes a subsequence of \mathbf{W} and θ denotes the model parameters.² Below, we describe how we use variations of this objective to finetune different LMs to generate questions, answers and distractors.³

Generating Synthetic Questions To train our question generator, we finetune the LM on the training question set $\{\mathbf{Q}^i\}$ to optimize

² $\mathbf{W}_{1:0}$ denotes an empty sequence

³Specific modifications for other tasks, e.g. textual entailment, are discussed in Appendix A.

¹<https://github.com/yangyiben/G-DAUG-c-Generative-Data-Augmentation-for-Commonsense-Reasoning>

the language modeling objective: $\mathcal{L}_q(\theta_q) = \sum_{i=1}^N \log P(\mathbf{Q}^i; \theta_q)$, where θ_q denotes the parameters of the question generator. After finetuning, we generate new questions with nucleus sampling (Holtzman et al., 2020), which is suitable for generating long-form text.

Generating Synthetic Answers and Distractors

To generate choice sets, we independently finetune two separate generative LMs, one for answers and the other for distractors. The answer and distractor generators are trained to maximize the conditional log-likelihood of the answer and the distractors, respectively, given the question. Mathematically, we optimize both $\mathcal{L}_a(\theta_a) = \sum_{i=1}^N \log P(\mathbf{C}_{y^i}^i | \mathbf{Q}^i; \theta_a)$ and $\mathcal{L}_d(\theta_d) = \sum_{i=1}^N \sum_{j \neq y^i} \log P(\mathbf{C}_j^i | \mathbf{Q}^i; \theta_d)$, where θ_a and θ_d denote the parameters of the answer and distractor generators, respectively. For answers, we use nucleus sampling with low temperature (for long answers) or greedy decoding (for short answers). To encourage diversity across generated distractors, we use nucleus sampling without temperature for these.

Data Relabeling. Our choice of generative LMs naturally defines labels for the synthetic choice sets. Alternatively, we consider using a supervised task model trained on the original training set, to *re-label* a candidate pool of synthetic answers and distractors. This is similar to treating the synthetic questions as unlabeled data and applying self-training. The utility of this self-training can be task-dependent; in our experiments, we used validation performance to determine whether or not to relabel our synthetic training data.

3 Synthetic Data Selection and Training

The above generation method can produce a large pool of examples, but training on all of them would be computationally expensive and might harm performance due to noisy generations. Here, we propose three data selection methods aimed at choosing more effective training examples from the generated pool (§3.1). Further, we outline a simple staged training procedure (§3.2) to mitigate the negative impact from noise in the synthetic data.

3.1 Selecting High-quality and Diverse Synthetic Examples

A randomly sampled synthetic dataset may contain examples that are similar to one another, along with low-quality generations (Holtzman et al., 2020).

We refer to such a random selection approach as **G-DAUG^c-Rand**. We hypothesize that a diverse and high-quality synthetic set would benefit the task model more. We present three data selection algorithms that target quality, diversity and a combination of both.

Filtering with Influence Functions. We hypothesize that filtering out detrimental synthetic training examples can boost downstream performance (Bras et al., 2020). A given training example x is considered *detrimental* if including x in the training set results in a higher generalization error, approximated by validation loss, i.e.:

$$\mathcal{L}(\mathcal{X}, \theta) = \frac{1}{|\mathcal{X}|} \sum_{x_i \in \mathcal{X}} l(x_i, \theta),$$

$$\mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{tr} \cup \{x\})) - \mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{tr})) > 0.$$

This would naively require retraining the model with x , which is computationally prohibitive. Fortunately, the validation loss change can be efficiently approximated through the use of influence functions (Atkinson et al., 1983; Koh and Liang, 2017). While previous work focuses on removing or perturbing existing training examples (Koh and Liang, 2017; Wang et al., 2018), we use influence functions to estimate the effect of including a novel synthetic example.

The main result from previous work (Atkinson et al., 1983; Koh and Liang, 2017) tells us that the influence of upweighting a training example x by some small ϵ on the model parameters $\hat{\theta}$ with the corresponding parameter space Θ is given by:

$$\hat{\theta}_{\epsilon, x} = \operatorname{argmin}_{\theta \in \Theta} \epsilon l(x, \theta) + \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i l(x_i, \theta)$$

$$\mathcal{I}_{up, params}(x) := \left. \frac{d\hat{\theta}_{\epsilon, x}}{d\epsilon} \right|_{\epsilon=0} = -H_{\hat{\theta}}^{-1} \nabla_{\theta} l(x, \hat{\theta}),$$

where w_i is weight for the training example x_i and $H_{\hat{\theta}}$ is the Hessian evaluated at $\hat{\theta}$. The above result is a slight generalization of Koh and Liang (2017), but it is straightforward to generalize their proof to the weighted empirical risk case. Then, we apply the chain rule to get the influence of upweighting x on the validation loss:

$$\mathcal{I}_{up, loss}(x) := \left. \frac{d\mathcal{L}(\mathcal{X}_{val}, \hat{\theta}_{\epsilon, x})}{d\epsilon} \right|_{\epsilon=0}$$

$$= \nabla_{\theta} \mathcal{L}(\mathcal{X}_{val}, \hat{\theta})^{\top} \mathcal{I}_{up, params}(x).$$

Note that $\mathcal{L}(\mathcal{X}_{tr}, \theta)$ can be rewritten as the following weighted average form to incorporate a new training example x_{new} :

$$\mathcal{L}(\mathcal{X}_{tr}, \theta) = \frac{1}{\sum_{i=1}^{N+1} w_i} \sum_{i=1}^{N+1} w_i l(x_i, \theta),$$

where $w_i = 1 \forall i \neq N+1$, $w_{N+1} = 0$ and $x_{N+1} = x_{new}$. Adding the new training example x_{new} is equivalent to upweighting x_{N+1} by $\frac{1}{N}$:

$$\begin{aligned} \mathcal{L}(\mathcal{X}_{tr} \cup \{x_{new}\}, \theta) &\propto \frac{1}{N} l(x_{N+1}, \theta) \\ &+ \frac{1}{\sum_{i=1}^{N+1} w_i} \sum_{i=1}^{N+1} w_i l(x_i, \theta). \end{aligned}$$

Applying the influence function $\mathcal{I}_{up,loss}(x)$, we obtain the following linear approximation of the validation loss change upon adding the training example x_{new} :

$$\begin{aligned} &\mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{tr} \cup \{x_{new}\})) - \mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{tr})) \\ &\approx \frac{1}{N} \mathcal{I}_{up,loss}(x_{new}). \end{aligned}$$

We adopt the stochastic estimation method described in Koh and Liang (2017) to efficiently compute $\mathcal{I}_{up,loss}$. Detrimental synthetic data will have $\frac{1}{N} \mathcal{I}_{up,loss} > 0$.

Another distinction between our approach and Koh and Liang (2017) is that they compute the influence of a single training example on a single *test* example, whereas we estimate influence of a synthetic training example on all validation examples at once, which makes our approach scalable to large pools of synthetic data. Our approach, referred to as **G-DAUG^c-Influence**, filters out detrimental synthetic data (i.e., the examples that have a positive estimated influence on the validation loss).

Selecting Diverse Examples. While G-DAUG^c-Influence promotes training data quality, it ignores diversity; we hypothesize that better diversity can provide a more reliable training signal. We propose a simple greedy algorithm that iteratively selects a synthetic training example from the pool that maximizes a diversity measure. Here, we use a simple measure of diversity equal to the number of unique unigrams in the selected training set. Surprisingly, preliminary experiments with a more sophisticated diversity method based on embedding distance did not improve results (see Appendix E for details).

We refer to this approach as **G-DAUG^c-Diversity** (see Algorithm 1).

Algorithm 1 G-DAUG^c-Diversity

Input: Synthetic data pool \mathcal{D}_{pool} , Target size N

Output: Synthetic dataset

Initialization: $\mathcal{D}_{synthetic} \leftarrow \{\}$

repeat

$$x_{max} = \operatorname{argmax}_{x \in \mathcal{D}_{pool}} \#n\text{-grams}(\mathcal{D}_{synthetic} \cup \{x\}) - \#n\text{-grams}(\mathcal{D}_{synthetic})$$

Add x_{max} to $\mathcal{D}_{synthetic}$

Remove x_{max} from \mathcal{D}_{pool}

until $|\mathcal{D}_{synthetic}| = N$

return $\mathcal{D}_{synthetic}$

Combining Influence Filtering and Diversity Maximization

G-DAUG^c-Influence and G-DAUG^c-Diversity have complementary benefits—the former aims at improving the quality of individual examples by filtering out detrimental ones, and the latter is designed to compose a diverse training set but does not consider quality. To reap both benefits, we propose a combined selection technique, **G-DAUG^c-Combo**, that first filters the data using G-DAUG^c-Influence, then selects examples according to G-DAUG^c-Diversity.

3.2 Training with Synthetic Data

In traditional data augmentation, new data is usually mixed with the original training examples to create an augmented training set (Wei and Zou, 2019; Kafle et al., 2017). However, when augmenting with data produced using a generative model, label noise can be detrimental to learning (Kafle et al., 2017). Moreover, the generated questions themselves can be noisy, i.e. nonsensical or ambiguous (see Table 7 under §4.2). To address this issue, we propose a simple training procedure that treats the synthetic and original data differently. We first train a model on the synthetic data (**Synthetic Training**), then further train on the original, human-authored training set (**Organic Training**). The motivation is to correct any unfavorable noise that may have been learnt during the first stage, by subsequently training on original data as more recent training data is favored by neural models (Goodfellow et al., 2014).

We also experiment with a mixing approach that minimizes a weighted average of the loss for the synthetic data and the original data, with an importance weight to downweight the synthetic examples

to mitigate noise. We find that two-stage training performs better than the importance-weighted loss (see Section 5).

4 Experiments

We present experiments on four commonsense multiple choice QA benchmarks: COMMONSENSEQA (Talmor et al., 2019), WINOGRANDE (Sakaguchi et al., 2020), CODAH (Chen et al., 2019) and HellaSwag (Zellers et al., 2019). Our techniques are also directly applicable to other closed-book multiple choice QA setups, such as science QA, and to textual entailment tasks with minor modifications. To evaluate G-DAUG^c's extensibility to these settings, we also experiment with a textual entailment task, SNLI (Bowman et al., 2015), and a closed-book version of the ARC-Challenge Scientific QA task (Clark et al., 2018) in which access to the scientific corpus for the ARC dataset (or any other information sources) is disallowed during test. We simulate low-resource settings on the large HellaSwag and SNLI datasets by downsampling these to 2K and 3K training samples respectively; the other data sets are either already low-resource or have a low-resource component. Dataset details are provided in Appendix A.

Robustness Evaluation In addition to measuring in-distribution performance, we also analyze robustness to perturbed or adversarial data. Following Wei and Zou (2019), we perform WordNet-based (Fellbaum, 1998) synonym replacement on the validation or test set (when test labels are available) with a 10% replacement rate.⁵ Our second evaluation with TextFooler (Jin et al., 2020) identifies the most important words and replaces these with the most semantically and grammatically correct substitutes, until the model prediction is altered. We adopt two metrics to measure robustness under TextFooler's attacks: 1) **failure rate**: the proportion of examples for which TextFooler fails to change the prediction and 2) **average perturbation ratio**: the average fraction of words replaced when TextFooler succeeds in altering a prediction. We re-implement TextFooler with two minor changes: we only swap words in questions, not answers, and we replace the Universal Sentence Encoder with SROBERTA (Reimers and Gurevych,

⁴<https://leaderboard.allenai.org/winogrande/submissions/public>, <https://www.tau-nlp.org/csqa-leaderboard>

⁵https://github.com/jasonwei20/eda_nlp

2019).

4.1 Experimental Settings

We use ROBERTA (Liu et al., 2019) as our pre-trained task model, and GPT-2 (Radford et al., 2019) as our pretrained generator.⁶ We use validation performance to decide whether to do relabeling for COMMONSENSEQA and WINOGRANDE, and apply relabeling by default on all other tasks (tuning this choice may boost performance). To perform a controlled comparison, we restrict the synthetic set size to be equal across all methods. We repeat all experiments with 10 random restarts and pick the best model based on validation performance. Additional experimental details, with hyperparameters, are provided in Appendix C.

Baselines Our first baseline is a finetuned ROBERTA model with no augmentation. We compare with existing work on data augmentation via a BACKTRANSLATION approach from Xie et al. (2019); under our setting the original and back-translated data are mixed at random.⁷

4.2 In-Distribution Results

Our main results for commonsense question answering are reported in Table 1. All G-DAUG^c variants outperform the baselines, highlighting the impact of generative data augmentation. On average, every other variant achieves higher test performance than G-DAUG^c-Rand, which further highlights the importance of our data selection approaches. In addition, influence and diversity selection methods score similarly, however, their combination (in G-DAUG^c-combo) outperforms either alone, which suggests that they are complementary selection approaches. More specifically, G-DAUG^c-Combo performs the best on 3/4 tasks and obtains the highest average score. Further, G-DAUG^c-Combo provides a 5.0% absolute gain over previously published state-of-the-art results on WINOGRANDE.⁸ For COMMONSENSEQA, G-DAUG^c-Combo outperforms the previous non-ensemble state-of-the-art (Zhu et al., 2020) by 0.4%. We also achieve a new state-of-the-art on CODAH, where the previous best (BERT-based) score was 67.5% (Chen et al., 2019). We find

⁶We used the HuggingFace library (Wolf et al., 2019).

⁷<https://github.com/google-research/uda/>

⁸These results are state-of-the-art for our model class; higher scores have been obtained using a T5 model with roughly an order of magnitude more parameters than ours.

	CSQA (Acc)	WINOGRANDE (AUC)	CODAH (Acc)	HellaSwag-2K (Acc)	Average
ROBERTA (reported)	72.1	66.4	-	-	-
ROBERTA (ours)	71.6	67.5	82.3	75.4	74.2
BACKTRANSLATION	70.2	67.2	81.8	73.0	73.1
G-DAUG ^c -Rand	71.8	70.9	83.6	75.9	75.6
G-DAUG ^c -Influence	72.1	70.9	84.3	75.8	75.8
G-DAUG ^c -Diversity	72.3	71.2	83.5	76.1	75.8
G-DAUG ^c -Combo	72.6	71.4	84.0	76.8	76.2

Table 1: Results on the test sets of four commonsense benchmarks. ROBERTA (reported) is the result for the ROBERTA-large baseline reported on public leaderboards.⁴ROBERTA (ours) is re-evaluation of the ROBERTA-large model using our setup. All G-DAUG^c methods outperform the baseline methods, and G-DAUG^c-Combo performs the best overall.

	CSQA	WINOGRANDE	CODAH	HellaSwag-2K	Average
ROBERTA (ours)	69.9	63.8	74.7	63.2	67.9
BACKTRANSLATION	69.0	62.3	75.5	65.4	68.1
G-DAUG ^c -Rand	72.1	65.5	75.9	64.1	69.4
G-DAUG ^c -Influence	71.0	65.7	76.2	64.3	69.3
G-DAUG ^c -Diversity	71.6	66.0	76.0	64.8	69.6
G-DAUG ^c -Combo	72.0	66.0	76.0	65.2	69.8

Table 2: Results on WordNet-based synonym replacement sets. For CODAH and HellaSwag-2K, we perturb test sets, as the labels are available. G-DAUG^c-Combo achieves the highest average score.

that BACKTRANSLATION hurts performance, and uniformly underperforms the ROBERTA baseline. See Appendix B for validation set results.

4.3 Robustness Results

Table 2 presents our evaluation on synonym replacement sets. The G-DAUG^c variants outperform the baselines, and G-DAUG^c-Combo obtains the best average performance. Table 3 shows results on the TextFooler adversarial attacks. Models trained with data augmentation are more robust to adversarial attacks, as all G-DAUG^c variants and BACKTRANSLATION outperform the ROBERTA baseline on both metrics. G-DAUG^c-Diversity obtains the best failure rate and average perturbation ratio (higher is better, in both metrics), and G-DAUG^c-Combo performs comparably with slightly worse numbers. Overall, the findings suggest that optimizing diversity increases robustness.

4.4 Results on ARC and SNLI

We explore G-DAUG^c's applicability outside of the commonsense domain in Table 4, via evaluation on the closed-book ARC-Challenge Scientific QA. Valid science questions are hard to generate because their semantics need to be precise, and we find that many of G-DAUG^c's generations for ARC are noisy. Perhaps surprisingly, nonetheless G-DAUG^c outperforms the baselines by a

large margin. G-DAUG^c-Influence achieves the best in-distribution performance, while G-DAUG^c-Diversity is the most robust against TextFooler but has worse accuracy than G-DAUG^c-Rand. This may suggest that optimizing for quality is more important when the synthetic data is noisier.

We also evaluate G-DAUG^c on a textual entailment using the SNLI dataset (Bowman et al., 2015) in Table 4. This task has a different format; it is a pair-wise classification task with 3 labels (details in Appendix A). We find that G-DAUG^c slightly improves accuracy and robustness over baselines. The performance is likely affected by a label skew introduced by influence-based filtering.

5 Analysis and Discussion

We now analyze G-DAUG^c's performance, focusing on WINOGRANDE where G-DAUG^c offers the most benefit. We first identify several factors that affect performance, and then present evidence that G-DAUG^c works by transferring knowledge from the pretrained generator to the task model.

5.1 Factors that Affect G-DAUG^c's Performance

G-DAUG^c is effective at different training sizes. Figure 3 illustrates that our winning strategy, G-DAUG^c-Combo, remains effective as the amount of training data varies, for WINOGRANDE.

	CSQA	WINOGRANDE	CODAH	Hellaswag-2K	Average
ROBERTA (ours)	14.8 / 12.6	4.5 / 7.8	30.9 / 15.8	17.4 / 9.8	16.9 / 11.5
BACKTRANSLATION	17.0 / 12.9	5.0 / 8.2	37.1 / 15.9	20.2 / 10.2	19.8 / 11.8
G-DAUG ^c -Rand	15.6 / 13.0	5.7 / 8.4	36.2 / 15.9	20.0 / 10.6	19.4 / 12.0
G-DAUG ^c -Influence	16.3 / 12.8	5.4 / 8.4	34.9 / 15.8	19.2 / 10.7	19.0 / 11.9
G-DAUG ^c -Diversity	16.0 / 12.9	5.9 / 8.4	36.1 / 16.2	21.4 / 10.4	19.9 / 12.0
G-DAUG ^c -Combo	16.5 / 12.6	5.9 / 8.5	35.2 / 15.7	21.3 / 10.5	19.7 / 11.8

Table 3: Robustness to TextFooler-based adversarial attacks (failure rate / average perturbation ratio, higher is better for both). Models trained with augmented data are more robust to TextFooler’s attacks compared to models without data augmentation. On average, G-DAUG^c-Diversity performs the best.

	ARC-Challenge Scientific QA					SNLI-3K					
	Val.	Test	Syn.	TF:Fail	TF:Pert	Val.	Test	Syn.	TF:Fail	TF:Pert	NLI Diag.
RoBERTa (ours)	43.5	39.4	35.2	6.6	9.3	91.8	88.6	77.5	17.0	20.2	56.7
Backtranslation	43.1	43.1	42.4	6.6	9.3	91.2	8.1	81.0	18.8	21.7	54.0
G-DAUG ^c -Rand	50.8	48.1	43.4	12.9	10.8	91.8	89.0	78.6	17.7	20.6	57.4
G-DAUG ^c -Influence	51.5	48.5	45.2	12.4	11.0	92.3	88.7	78.6	18.0	20.7	56.9
G-DAUG ^c -Diversity	49.5	47.5	42.2	13.9	10.8	92.0	89.0	79.4	19.0	20.5	57.7
G-DAUG ^c -Combo	50.8	48.2	43.8	13.1	10.7	91.9	88.7	78.7	16.7	20.5	57.6

Table 4: Results on closed-book ARC-Challenge Scientific QA and SNLI-3K, along with robustness to synonym replacement, TextFooler (TF) attacks and NLI Diagnostics. G-DAUG^c improves accuracy and robustness.

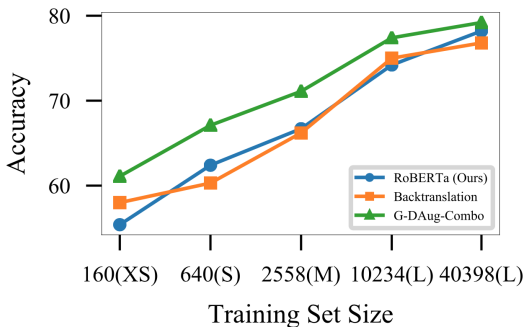


Figure 3: Validation results for different training set sizes on the WINOGRANDE dataset (in log scale). G-DAUG^c helps more for smaller training sizes.

The improvement over baseline is largest in the low-resource (small training size) regime. For the smallest sizes, XS and S, G-DAUG^c-Combo increases the *effective training size* by a factor of 4 (i.e. training on XS or S matches unaugmented ROBERTA’s performance on S or M, respectively). In contrast, BACKTRANSLATION only helps for the XS size, but hurts performance on larger sizes.

Staged training is essential. G-DAUG^c uses a two-staged training method (Section 3.2) aimed at mitigating the effect of noise in the generated data. We analyze alternative training protocols on the WINOGRANDE-L dataset: *Mixing* (training on the union of generated and original data) and Importance Weighted Loss. Compared to a no-augmentation baseline (with accuracy of 75.9),

two stage training (+1.8 increase) outperforms both mixing (+0.0) and importance weighted loss (+0.7).

	Random	Influence	Diversity	Whole Pool
Size	127478	127478	127478	380700
Acc	71.7	74.4	73.0	73.1

Table 5: Results comparing G-DAUG^c’s filtering methods against using the entire synthetic data pool for augmentation, on WINOGRANDE-M.

Filtering synthetic data does not hurt accuracy. G-DAUG^c’s filtering methods are designed to identify a high-quality and diverse subset of the generated data, to reduce training cost (compared to training on the entire generated pool) without harming accuracy. We evaluate whether G-DAUG^c is successful at achieving this in Table 5, by comparing G-DAUG^c against using the entire synthetic data pool for G-DAUG^c-Influence and G-DAUG^c-Diversity.⁹ The selection approaches provide comparable or better accuracy compared to using the entire pool, despite using three times less data.

5.2 Why Does G-DAUG^c Work?

Below, we present analysis suggesting that G-DAUG^c works by transferring knowledge from the pretrained model to the task model. In particular, we find that using a pre-trained generator is

⁹G-DAUG^c-Combo utilizes a larger pool, so it is not comparable.

critical, and that the generated questions are often coherent, include new semantic units, and carry informative labels.

Using a Pretrained Generator is critical. We analyze the impact of the pretrained generator by comparing our standard G-DAUG^c-Rand setting with a setting where the generator is *not* pretrained, but instead trained from scratch. We find that using GPT-2 trained from scratch results in a score of 67.8% on the WINOGRANDE-M validation set. This is a slight improvement (by 0.2%) over the unaugmented baseline, but is far inferior to the 3.9% improvement obtained when using the pretrained GPT-2. This suggests that using a pretrained generator is critical for G-DAUG^c.

	WINOGRANDE-L	CSQA
Baseline	75.9	77.1
Generator label	76.2	78.1
Random relabeling	66.8	77.1
Model relabeling	77.7	77.7

Table 6: Validation accuracy of G-DAUG^c with different labeling methods on WINOGRANDE-L and COMMONSENSEQA. Random labels hurt accuracy, and model relabeling helps on WINOGRANDE but not on COMMONSENSEQA.

Synthetic data labels are important. Even fully unsupervised language model pretraining can improve performance, when using task-relevant data (Gururangan et al., 2020). This raises the question of whether G-DAUG^c boosts performance by simply exposing the model to more task-relevant text, or if the generated labels are in fact informative. A related question is whether G-DAUG^c’s optional self-supervised relabeling improves performance. We analyze these questions for WINOGRANDE-L and COMMONSENSEQA in Table 6, evaluating G-DAUG^c with three labeling methods: (i) generator labels, (ii) random relabeling, and (iii) relabeling with a task model. When the generator labels are flipped randomly, G-DAUG^c is unable to outperform the baselines for either dataset (in fact, it dramatically underperforms on WINOGRANDE-L). This implies that the correctness of labels is crucial for G-DAUG^c. Self-supervised relabeling provides a 1.5% absolute gain in WINOGRANDE-L, but a 0.4% drop in COMMONSENSEQA, which suggests its utility is task-dependent.

G-DAUG^c introduces new semantic units. We investigate *how distinct* the generated questions

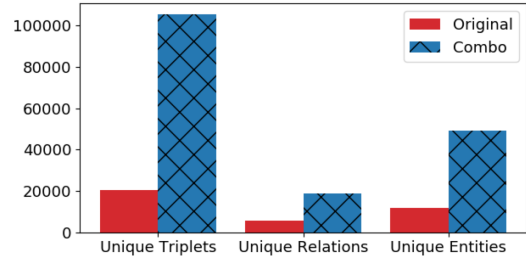


Figure 4: OpenIE analysis on the original data and synthetic data used by G-DAUG^c-Combo on WINOGRANDE-M. The synthetic dataset contains many more unique semantic units compared to the original dataset.

are from each other and from the original training data. We observe that G-DAUG^c only rarely generates exact duplicate questions (e.g., on COMMONSENSEQA, 0.06% of the questions are duplicates). We further investigate if G-DAUG^c introduces new entities and relations to the training data, or if it merely reuses the ones found in the original training set. We quantify the diversity of our synthetic dataset compared to the original data by counting the number of unique semantic units produced by performing Open Information Extraction (Banko et al., 2007) on the data. Specifically, we run the Stanford Open IE package (Angeli et al., 2015) and report the number of unique triplets, relations and entities extracted from our WINOGRANDE-M datasets in Figure 4. The synthetic data includes many more unique semantic units than the original training data, suggesting that G-DAUG^c does introduce new semantic units in the training set.

G-DAUG^c produces mostly fluent questions.

To evaluate G-DAUG^c’s output for fluency, we employ three human annotators to rate generated COMMONSENSEQA questions for their coherence and answerability on a scale of 1 to 4, where a rating of 3 denotes an acceptable question. We obtained a total of 1,387 labels. We measured annotator agreement on a separate set of 50 questions, obtaining a Fleiss’ Kappa of 0.41, which is at the low end of moderate annotator agreement, acceptable given the subjective nature of the task. A large (74.04%) majority of questions met the acceptability threshold, with an overall average rating of 3.34. Examples are shown in Table 7.

Next, we ask annotators to answer the 1,027 acceptable questions, where they can edit choices (but not questions) if they are unable to pick a unique correct answer from the given choices. The

Rating	Description	Examples	Count	Pct.
1	Nonsensical	What is a square leg made of made out of? What country does a cow go to make a milk run?	54	3.89%
2	Ambiguous or unanswerable	A person is a human, but they are called what? He hated flying, the controls were what?	306	22.06%
3	Minor errors (e.g., grammar)	What do you put on your head to do when you're swimming? Where does a bugle call be played?	138	9.95%
4	Coherent and Fluent	What is a person likely to feel when applying for jobs? If you're running late for work what would you be doing?	889	64.10%

Table 7: Examples and prevalence of generated commonsense questions with different manually-assigned fluency ratings, for the COMMONSENSEQA dataset. Ratings of 3 and higher correspond to questions that are answerable and address common sense, and most of G-DAUG^c's generated questions fall into this category.

editing rate is relatively high, at 55.3%. We mix these human-labeled examples with the original training set to train a ROBERTA model, and obtain 78.1% validation accuracy, which is comparable to G-DAUG^c, despite using approximately 50x fewer questions. This suggests that human labels can provide higher leverage than the noisy labels from G-DAUG^c, although human labeling is expensive.

Additional analyses, provided in Appendix F, show that model sharpness approximated by the Hessian trace (Yao et al., 2019) does not completely explain G-DAUG^c's performance; and, G-DAUG^c is more effective than ensembling with a finetuned generator.

6 Related Work

Data augmentation is a common practice in computer vision, where it takes the form of image transformations like translation and rotation (Perez and Wang, 2017). For language tasks, data augmentation is less straightforward. Broadly, previous augmentation methods have used back-translation architectures (Sennrich et al., 2016; Xie et al., 2019), heuristics based on syntactic and semantic properties of text including word replacements using a thesaurus (Zhang et al., 2015; Wei and Zou, 2019) and word embeddings (Wang and Yang, 2015; Fadaee et al., 2017; Kobayashi, 2018; Wu et al., 2019), and recently, generative models for synthesizing novel examples for text classification and reading comprehension (Anaby-Tavor et al., 2020; Kumar et al., 2020; Puri et al., 2020b). Our framework is similar to the last of these as we focus on generative models for data augmentation, but our work is the first to present a generative approach for the challenging commonsense QA setting, and we introduce new data selection approaches to improve the informativeness and diversity of synthetic data.

Concurrently, there has been work on generat-

ing adversarial examples for analyzing black-box classifiers. These approaches use generative adversarial networks (Zhao et al., 2018b) and population-based optimization algorithms (Alzantot et al., 2018). Previous work has also presented methods to generate questions for reading comprehension (Heilman and Smith, 2010; Rus et al., 2011; Alberti et al., 2019; Puri et al., 2020a), online tutoring (Lindberg et al., 2013), factual QA (Serban et al., 2016) and visual question generation (Mostafazadeh et al., 2016). A comprehensive survey on neural question generation can be found in Pan et al. (2019). Our work is distinct in that it targets question generation in a closed-book setting, investigates the generation of *answers* as well as *distractors*, and is aimed at data augmentation.

7 Conclusion

We introduced G-DAUG^c, a novel data augmentation framework to generate synthetic training data, preserving quality and diversity. We demonstrate that G-DAUG^c is effective on multiple commonsense reasoning benchmarks, with improvements on in-distribution performance, as well as robustness against perturbed evaluation sets and challenge sets. Our analysis shows that G-DAUG^c tends to perform better in low-resource settings and that our data selection strategies are important for performance. Future work might explore more sophisticated methods to enhance quality and diversity of generated training data, including having humans-in-the-loop for relabeling.

Acknowledgments

This work was supported in part by NSF Grant IIS-1351029. We thank Iz Beltagy, Jonathan Bragg, Isabel Cachola, Arman Cohan, Mike D'Arcy, Daniel King, Kyle Lo, and Lucy Lu Wang for helpful comments.

References

- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. [Don't just assume; look and answer: Overcoming priors for visual question answering](#). *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Chris Alberti, Daniel Andor, Emily Pitler, Jacob Devlin, and Michael Collins. 2019. [Synthetic QA corpora generation with roundtrip consistency](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6168–6173, Florence, Italy. Association for Computational Linguistics.
- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. [Generating natural language adversarial examples](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.
- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. [Not enough data? Deep learning to the rescue!](#) In *Proc. of AAAI*.
- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. [Leveraging linguistic structure for open domain information extraction](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 344–354, Beijing, China. Association for Computational Linguistics.
- A. C. Atkinson, R. Dennis Cook, and Sanford Weisberg. 1983. [Residuals and influence in regression](#).
- Haim Avron and Sivan Toledo. 2011. [Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix](#). *J. ACM*, 58:8:1–8:34.
- Michele Banko, Michael J. Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. [Open information extraction from the web](#). In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, pages 2670–2676, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Yonatan Belinkov and Yonatan Bisk. 2017. [Synthetic and natural noise both break neural machine translation](#). *arXiv preprint arXiv:1711.02173*.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Çelikyilmaz, and Yejin Choi. 2019. [Comet: Commonsense transformers for automatic knowledge graph construction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. [A large annotated corpus for learning natural language inference](#). *arXiv preprint arXiv:1508.05326*.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew E. Peters, Ashish Sabharwal, and Yejin Choi. 2020. [Adversarial filters of dataset biases](#).
- Michael Chen, Mike D'Arcy, Alisa Liu, Jared Fernandez, and Doug Downey. 2019. [CODAH: An adversarially-authored question answering dataset for common sense](#). In *Proceedings of the 3rd Workshop on Evaluating Vector Space Representations for NLP*, pages 63–69, Minneapolis, USA. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *ArXiv*, abs/1803.05457.
- Joe Davison, Joshua Feldman, and Alexander Rush. 2019. [Commonsense knowledge mining from pre-trained models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1173–1178, Hong Kong, China. Association for Computational Linguistics.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. [Learning to ask: Neural question generation for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1342–1352, Vancouver, Canada. Association for Computational Linguistics.
- Marzieh Fadaee, Arianna Bisazza, and Christof Monz. 2017. [Data augmentation for low-resource neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 567–573, Vancouver, Canada. Association for Computational Linguistics.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Maxwell Forbes and Yejin Choi. 2017. [Verb physics: Relative physical knowledge of actions and objects](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1:*

- Long Papers*), pages 266–276, Vancouver, Canada. Association for Computational Linguistics.
- Ian J. Goodfellow, Mehdi Mirza, Xia Da, Aaron C. Courville, and Yoshua Bengio. 2014. [An empirical investigation of catastrophic forgetting in gradient-based neural networks](#). *CoRR*, abs/1312.6211.
- Jonathan Gordon and Benjamin Van Durme. 2013. [Reporting bias and knowledge acquisition](#). In *Proceedings of the 2013 workshop on Automated knowledge base construction*, pages 25–30.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of ACL*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Yaru Hao, Li Dong, Furu Wei, and Ke Xu. 2019. [Visualizing and understanding the effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4143–4152, Hong Kong, China. Association for Computational Linguistics.
- Michael Heilman and Noah A. Smith. 2010. [Good question! statistical ranking for question generation](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 609–617, Los Angeles, California. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Flat minima](#). *Neural Computation*, 9:1–42.
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). *International Conference on Learning Representations*.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. [Adversarial example generation with syntactically controlled paraphrase networks](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1875–1885, New Orleans, Louisiana. Association for Computational Linguistics.
- Robin Jia and Percy Liang. 2017. [Adversarial examples for evaluating reading comprehension systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. [Is BERT really robust? natural language attack on text classification and entailment](#).
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. [Billion-scale similarity search with gpus](#). *arXiv preprint arXiv:1702.08734*.
- Kushal Kafle, Mohammed Yousefhussien, and Christopher Kanan. 2017. [Data augmentation for visual question answering](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 198–202, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2016. [On large-batch training for deep learning: Generalization gap and sharp minima](#). *ArXiv*, abs/1609.04836.
- Sosuke Kobayashi. 2018. [Contextual augmentation: Data augmentation by words with paradigmatic relations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 452–457, New Orleans, Louisiana. Association for Computational Linguistics.
- Pang Wei Koh and Percy Liang. 2017. [Understanding black-box predictions via influence functions](#). In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 1885–1894. PMLR.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2020. [Data augmentation using pre-trained transformer models](#). *ArXiv*, abs/2003.02245.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2011. [The winograd schema challenge](#). In *KR*.
- David Lindberg, Fred Popowich, John Nesbit, and Phil Winne. 2013. [Generating natural language questions to support learning on-line](#). In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 105–114, Sofia, Bulgaria. Association for Computational Linguistics.
- 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](#). *arXiv preprint arXiv:1907.11692*.
- Shangwen Lv, Daya Guo, Jingjing Xu, Duyu Tang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, and Songlin Hu. 2020. [Graph-based](#)

- reasoning over heterogeneous external knowledge for commonsense question answering. In *Proc. of AAAI*.
- Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Margaret Mitchell, Xiaodong He, and Lucy Vanderwende. 2016. [Generating natural questions about an image](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1802–1813, Berlin, Germany. Association for Computational Linguistics.
- Liangming Pan, Wenqiang Lei, Tat-Seng Chua, and Min-Yen Kan. 2019. [Recent advances in neural question generation](#). *arXiv preprint arXiv:1905.08949*.
- Luis Perez and Jason Wang. 2017. [The effectiveness of data augmentation in image classification using deep learning](#). *arXiv preprint arXiv:1712.04621*.
- R. Puri, Ryan Spring, M. Patwary, M. Shoeybi, and Bryan Catanzaro. 2020a. Training question answering models from synthetic data. *ArXiv*, abs/2002.09599.
- Raul Puri, Ryan Spring, Mostofa Patwary, Mohammad Shoeybi, and Bryan Catanzaro. 2020b. [Training question answering models from synthetic data](#). *arXiv preprint arXiv:2002.09599*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *arXiv e-prints*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Vasile Rus, Brendan Wyse, Paul Piwek, Mihai Lintean, Svetlana Stoyanchev, and Cristian Moldovan. 2011. [Question generation shared task and evaluation challenge – status report](#). In *Proceedings of the 13th European Workshop on Natural Language Generation*. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [WINOGRANDE: An adversarial winograd schema challenge at scale](#). In *Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, New York, USA.
- Roy Schwartz, Maarten Sap, Ioannis Konstas, Leila Zilles, Yejin Choi, and Noah A. Smith. 2017. [The effect of different writing tasks on linguistic style: A case study of the ROC story cloze task](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 15–25, Vancouver, Canada. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Iulian Vlad Serban, Alberto García-Durán, Caglar Gulcehre, Sungjin Ahn, Sarath Chandar, Aaron Courville, and Yoshua Bengio. 2016. [Generating factoid questions with recurrent neural networks: The 30M factoid question-answer corpus](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 588–598, Berlin, Germany. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Trieu H. Trinh and Quoc V. Le. 2018. [A simple method for commonsense reasoning](#). *ArXiv*, abs/1806.02847.
- Tianyang Wang, Jun Huan, and Bo Li. 2018. [Data dropout: Optimizing training data for convolutional neural networks](#). *2018 IEEE 30th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 39–46.
- William Yang Wang and Diyi Yang. 2015. [That’s so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2557–2563, Lisbon, Portugal. Association for Computational Linguistics.
- Jason Wei and Kai Zou. 2019. [EDA: Easy data augmentation techniques for boosting performance on text classification tasks](#). In *Proc. of EMNLP-IJCNLP*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s trans-](#)

- formers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Xing Wu, Shangwen Lv, Liangjun Zang, Jizhong Han, and Songlin Hu. 2019. [Conditional bert contextual augmentation](#). In *International Conference on Computational Science*, pages 84–95. Springer.
- Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. 2019. [Unsupervised data augmentation for consistency training](#). *arXiv preprint arXiv:1904.12848*.
- Yiben Yang, Larry Birnbaum, Ji-Ping Wang, and Doug Downey. 2018. [Extracting commonsense properties from embeddings with limited human guidance](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 644–649, Melbourne, Australia. Association for Computational Linguistics.
- Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael W. Mahoney. 2019. [Pyhessian: Neural networks through the lens of the hessian](#). *ArXiv*, abs/1912.07145.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A large-scale adversarial dataset for grounded commonsense inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, pages 649–657.
- Yao Zhao, Xiaochuan Ni, Yuanyuan Ding, and Qifa Ke. 2018a. [Paragraph-level neural question generation with maxout pointer and gated self-attention networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3901–3910, Brussels, Belgium. Association for Computational Linguistics.
- Zhengli Zhao, Dheeru Dua, and Sameer Singh. 2018b. [Generating natural adversarial examples](#). In *International Conference on Learning Representations*.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. [Neural question generation from text: A preliminary study](#). *arXiv preprint arXiv:1704.01792*.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jing jing Liu. 2020. [FreeLb: Enhanced adversarial training for natural language understanding](#). In *ICLR 2020*.

A Datasets

CommonsenseQA (Talmor et al., 2019): CommonsenseQA is a multiple choice QA dataset that consists of 12,247 examples, which aims to test commonsense reasoning capabilities. We use the official random split 1.11 which is an 80/10/10 split. We apply greedy decoding to generate answers, as answers are fairly short for this dataset.

WINOGRANDE (Sakaguchi et al., 2020): WINOGRANDE is a benchmark for commonsense reasoning, inspired by the original Winograd Schema Challenge design (Levesque et al., 2011), with a larger dataset size and higher difficulty level. It consists of 44K questions with five different training sizes: 160, 640, 2,558, 10,234 and 40,398 questions. The evaluation metric is Area Under the (learning) Curve. We observe that applying top-2 greedy decoding on the answer generator is able to yield a satisfactory set of choices, so the distractor generator is not used in this task. The Winograd schema requires that questions in twin pairs have opposite labels (Levesque et al., 2011). We use the following method to generate twin questions: 1. generate a sequence until a blank symbol ”_” is produced. 2. use two independent runs of sampling to complete the question in two different ways to form twins. The above process does not guarantee that the labels will differ for the two twins, so we further filter out generated pairs that do not have different labels.

CODAH (Chen et al., 2019): CODAH is an adversarially-constructed benchmark which tests commonsense reasoning using sentence-completion questions, inspired by the Swag dataset (Zellers et al., 2018). It contains 2,801 questions in total, and uses 5-fold cross validation for evaluation.¹⁰ We lower the temperature to 0.5 for the answer generation in order to increase the confidence of the generated answers.

HellaSwag (Zellers et al., 2019): HellaSwag is a more challenging version of the Swag dataset (Zellers et al., 2018), and the task is similar to CODAH. The dataset consists of 70K questions where each question comes from one of two domains: ActivityNet or WikiHow. In order to test our methods under a low-resource setting, we downsample the training set to 2,000 examples. We take a random

¹⁰The original CODAH work does not specify a particular 5-fold split, so we choose these randomly. We will release our splits for replicability.

sample of 1000 questions from the original validation set to serve as our validation data, and another non-overlapping random sample of 5,000 questions from the same set as our test data. The generation settings are the same as CODAH’s.

SNLI (Bowman et al., 2015): SNLI is a natural language inference dataset with 570K pairs of labeled sentences. The label assigned to each sentence pair is one of entailment, contradiction or neutral. For low-resource experiments, we downsample the dataset to 3K training examples, which contains 1K unique premises and a hypothesis for all three labels. Similarly, we use a downsampled development set with 999 examples (333 premises and 3 hypotheses for each label). The generative model is fine-tuned by providing the premise, label and hypothesis, separated by special delimiters marking the beginning and end of each element.

ARC-Challenge (Clark et al., 2018): The ARC Dataset consists of 7,787 natural grade-school science questions that are used on standardized tests. The ARC-Challenge Set contains 2,590 questions answered incorrectly by both a retrieval-based algorithm and a word co-occurrence algorithm. We use the official split, which has 1,119 train, 299 validation, and 1,172 test examples. The generation settings are the same as COMMONSENSEQA’s.

B Validation Set Results

In Table 8, we summarize our main results on the validation sets, comparing the G-DAUG^c methods against an unaugmented baseline and a backtranslation augmentation baseline. All G-DAUG^c methods consistently outperform the baseline methods in every benchmark. The proposed selection methods provide an extra boost on average, compared to G-DAUG^c-Rand. Among those, G-DAUG^c-Influence achieves the best performance across all tasks, which is expected as G-DAUG^c-Influence selects examples which are helpful in reducing validation loss. Interestingly, G-DAUG^c-Combo scores lower than G-DAUG^c-Influence, although it outperforms G-DAUG^c-Diversity. Finally, backtranslation does not demonstrate any benefit and obtains lower results compared to the augmented baseline in all benchmarks.

Method	CSQA (Acc)	WINOGRANDE (AUC)	CODAH (Acc)	HellaSwag-2K (Acc)	Average
ROBERTA (reported)	78.4	66.6	-	-	-
ROBERTA (ours)	77.1	68.4	84.2	75.2	76.2
Backtranslation	76.4	67.7	83.4	74.2	75.4
G-DAUG ^c -Rand	78.1	72.0	85.7	77.2	78.3
G-DAUG ^c -Influence	78.8	73.0	87.2	78.3	79.3
G-DAUG ^c -Diversity	78.1	72.8	86.0	76.6	78.4
G-DAUG ^c -Combo	78.2	72.7	86.7	77.5	78.8

Table 8: Results on the validation sets of four commonsense benchmarks. All G-DAUG^c methods outperform the baseline methods, in particular, G-DAUG^c-Influence performs the best on all tasks, which is expected as it selects examples which are helpful in reducing validation loss.

C Hyperparameter Settings and Input Formats

Hyperparameter settings for finetuning GPT-2, ROBERTA and G-DAUG^c are shown in Tables 11, 12, 14, 15 and 16. We manually tune the learning rate and the number of epochs for GPT-2 finetuning based on validation perplexity. For finetuning ROBERTA baseline models, we select the number of epochs from $\{1,3,5,8,10\}$ based on validation accuracy for CSQA, WINOGRANDE and HellaSwag-2K. For CODAH, SNLI-3K and ARC-Challenge, we simply use 5 epochs. For G-DAUG^c synthetic training, we train all models using a learning rate of $5e-6$ for one epoch. For G-DAUG^c organic training, we use the same hyperparameter settings as ROBERTA baselines (except for CSQA and HellaSwag-2K, where we find reducing 2 epochs gives significantly better results). In Tables 9 and 10, we specify the input formats for finetuning GPT-2 and ROBERTA. Finally, we benchmark the running time of our implementations of the influence and diversity selection methods on the task of selecting 127,478 examples from a pool consisting of 380,700 candidates for WINOGRANDE-M. We use one Nvidia 2080 Ti GPU and one Intel Core I9-7900X with 10 cores and a clockspeed of 3.3 GHz. The running time of the influence and diversity algorithms is about 8.3 hours and 2.9 hours, respectively.

D Influence Functions

In practice, since the generalization error is usually approximated by validation loss, a training example x_i is considered detrimental if it increases

validation loss, i.e.:

$$\mathcal{L}(\mathcal{X}, \theta) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} l(x, \theta), \quad (1)$$

$$\mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{train} \cup \{x_i\})) - \mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{train})) > 0, \quad (2)$$

where $\mathcal{X}_{train} = \{x_i\}_{i=1}^N$ is a training set, $\mathcal{X}_{val} = \{x_i\}_{i=1}^M$ is a validation set, l is a loss function, and $\hat{\theta}(\mathcal{X}_{train}) = \operatorname{argmin}_{\theta \in \Theta} \mathcal{L}(\mathcal{X}_{train}, \theta)$ is an empirical risk minimizer.

The main result from previous work (Atkinson et al., 1983; Koh and Liang, 2017) tells us that the influence of upweighting a training example x by some small ϵ on the model parameters $\hat{\theta}$ with the corresponding parameter space Θ is given by:

$$\hat{\theta}_{\epsilon, x} = \operatorname{argmin}_{\theta \in \Theta} \epsilon l(x, \theta) + \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i l(x_i, \theta) \quad (3)$$

$$\mathcal{I}_{up, params}(x) := \left. \frac{d\hat{\theta}_{\epsilon, x}}{d\epsilon} \right|_{\epsilon=0} = -H_{\hat{\theta}}^{-1} \nabla_{\theta} l(x, \hat{\theta}), \quad (4)$$

where w_i is weight for the training example x_i and $H_{\hat{\theta}} = \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i \nabla_{\theta}^2 l(x_i, \hat{\theta})$ is the Hessian evaluated at $\hat{\theta}$. The above result is a slight generalization of Koh and Liang (2017), since the simple average used in that work is a special case of our weighted average, but it is straightforward to generalize their proof to our weighted empirical risk case and we omit the details of the proof in this paper. Then, we apply the chain rule to get the influence of upweighting x on the validation loss:

$$\mathcal{I}_{up, loss}(x) := \left. \frac{d\mathcal{L}(\mathcal{X}_{val}, \hat{\theta}_{\epsilon, x})}{d\epsilon} \right|_{\epsilon=0} \quad (5)$$

$$= \nabla_{\theta} \mathcal{L}(\mathcal{X}_{val}, \hat{\theta})^{\top} \mathcal{I}_{up, params}(x). \quad (6)$$

Task	Format
CSQA	Q: Where can I stand on a river to see water falling without getting wet? A: waterfall $\langle /s \rangle$
WINOGRANDE	$\langle /s \rangle$ Feeling a draft, William asked Neil to please close the front door because _ was closer. $\langle /s \rangle$ Neil $\langle /s \rangle$
CODAH	$\langle /s \rangle$ I am always very hungry before I go to bed. I am $\langle /s \rangle$ concerned that this is an illness. $\langle /s \rangle$
HellaSwag-2K	$\langle /s \rangle$ A man is on a sandy beach, playing croquette. he $\langle /s \rangle$ is parasailing, making a random move. $\langle /s \rangle$
SNLI-3K	$\langle \text{PREM} \rangle$ Five black dogs run in a field. $\langle / \text{PREM} \rangle \langle \text{ANS} \rangle$ entailment $\langle / \text{ANS} \rangle \langle \text{HYP} \rangle$ Some animals running. $\langle / \text{HYP} \rangle$
ARC-Challenge	Q: Which of the following is an example of a physical change? A: breaking a glass $\langle /s \rangle$

Table 9: Input formats for GPT-2. "Q:" and "A:" are the prefix for a question and a candidate answer (choice).

Task	Format
CSQA	$\langle s \rangle$ Q: Where can I stand on a river to see water falling without getting wet? $\langle /s \rangle$ A: waterfall $\langle /s \rangle$
WINOGRANDE	$\langle s \rangle$ Feeling a draft, William asked Neil to please close the front door because _ was closer. $\langle /s \rangle$ Neil $\langle /s \rangle$
CODAH	$\langle s \rangle$ I am always very hungry before I go to bed. I am $\langle /s \rangle$ concerned that this is an illness. $\langle /s \rangle$
HellaSwag-2K	$\langle s \rangle$ A man is on a sandy beach, playing croquette. he $\langle /s \rangle$ is parasailing, making a random move. $\langle /s \rangle$
SNLI-3K	$\langle s \rangle$ Five black dogs run in a field. $\langle /s \rangle$ Some animals running. $\langle /s \rangle$
ARC-Challenge	$\langle s \rangle$ Q: Which of the following is an example of a physical change? $\langle /s \rangle$ A: breaking a glass $\langle /s \rangle$

Table 10: Input formats for ROBERTA. "Q:" and "A:" are the prefix for a question and a candidate answer (choice).

Note that $\mathcal{L}(\mathcal{X}_{train}, \theta)$ can be rewritten as the following weighted average form to incorporate a new training example x_{new} :

$$\mathcal{L}(\mathcal{X}_{train}, \theta) = \frac{1}{\sum_{i=1}^{N+1} w_i} \sum_{i=1}^{N+1} w_i l(x_i, \theta),$$

where $w_i = 1 \forall i \neq N+1$, $w_{N+1} = 0$ and $x_{N+1} = x_{new}$. Adding the new training example x_{new} is equivalent to upweighting x_{N+1} by $\frac{1}{N}$:

$$\begin{aligned} \mathcal{L}(\mathcal{X}_{train} \cup \{x_{new}\}, \theta) &= \frac{N}{N+1} \left(\frac{1}{N} l(x_{N+1}, \theta) \right) \\ &+ \frac{1}{\sum_{i=1}^{N+1} w_i} \sum_{i=1}^{N+1} w_i l(x_i, \theta) \\ &\propto \frac{1}{N} l(x_{N+1}, \theta) + \frac{1}{\sum_{i=1}^{N+1} w_i} \sum_{i=1}^{N+1} w_i l(x_i, \theta). \end{aligned}$$

Applying the influence function $\mathcal{I}_{up,loss}(x)$, we obtain the following linear approximation of the validation loss change upon adding the training example x_{new} :

$$\mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{train} \cup \{x_{new}\})) - \mathcal{L}(\mathcal{X}_{val}, \hat{\theta}(\mathcal{X}_{train})) \quad (7)$$

$$\approx \frac{1}{N} \mathcal{I}_{up,loss}(x_{new}). \quad (8)$$

We adopt the stochastic estimation method described in Koh and Liang (2017) to efficiently compute $\mathcal{I}_{up,loss}$. Detrimental synthetic data will have $\frac{1}{N} \mathcal{I}_{up,loss} > 0$.

E Diversity Selection using Embedding Distance

We define our embedding distance based diversity measure as the sum of the cosine distances between every pair of selected examples. To attempt to maximize this measure, we use a greedy algorithm that at each iteration randomly samples 10K candidate examples from the pool, and selects the candidate that maximizes the distance between it and its nearest neighbor in the set of examples selected so far. We use SROBERTA (Reimers and Gurevych, 2019) as our sentence embedding method and Faiss (Johnson et al., 2017) as our nearest neighbor searcher. We compare the embedding distance based measure with the unigram approach on WINOGRANDE dataset. The embedding distance based diversity selection is not found to be more effective than the unigram approach, in fact it performs 0.6% worse.

F Additional Analysis

Sharpness Analysis. Previous work (Hochreiter and Schmidhuber, 1997; Keskar et al., 2016; Yao et al., 2019) has shown that models with flatter local minima tend to generalize better. Moreover, Hao et al. (2019) show that pretraining helps BERT to achieve flat and wide optima in the finetuning stage, which partially explains its performance benefits. We investigate whether G-DAUG^c's data augmentation may also encourage flatter optima. Specifically, using the fact that a larger Hessian trace for a model implies a sharper local minimum (Yao et al., 2019), we compute the Hessian trace of 10 baseline and 10 G-DAUG^c-Combo methods using

Hyperparam	CSQA	WINOGRANDE	CODAH	HellaSwag-2K	SNLI-3K	ARC-Challenge
Version	Large	Medium	Medium	Medium	Large	Medium
Hardware	I9-7900X	RTX 2080Ti	RTX 2080Ti	RTX 2080Ti	RTX 8000	RTX 2080Ti
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW
Adam β_1	0.9	0.9	0.9	0.9	0.9	0.9
Adam β_2	0.98	0.98	0.98	0.98	0.999	0.98
Adam ϵ	1e-6	1e-6	1e-6	1e-6	1e-8	1e-6
Mixed Precision	No	Yes	Yes	Yes	Yes	Yes
LR (q/a/d)	1e-5/5e-6/2e-5	*	4e-5/5e-5/5e-5	4e-5/5e-5/5e-5	5e-5	2e-5/1e-5/1e-5
Epochs (q/a/d)	3/5/3	*	3/3/3	3/3/3	3	3/5/5
Grad Clipping	1.0	1.0	1.0	1.0	1.0	1.0
Weight Decay	0.01	0.01	0.01	0.01	0.0	0.01
Batch Size	16	16	16	16	16	16
Max Length (q/a/d)	62/70/70	72/72/-	62/92/92	62/128/128	128	90/120
Warmup Ratio	0.06	0.06	0.06	0.06	0.06	0.06
LR Decay	Linear	Linear	Linear	Linear	Linear	Linear

Table 11: Hyperparameter settings for finetuning GPT-2. "q/a/d" stands for "question/answer/distractor". Some hyperparameters for WINOGRANDE is shown in a separate table as they vary with the train size.

Hyperparam	XS	S	M	L	XL
LR (q/a)	5e-5/5e-5	2e-5/5e-5	2e-5/5e-5	2e-5/5e-5	1e-5/5e-5
Epochs (q/a)	8/12	6/6	3/3	3/3	3/1

Table 12: Hyperparameter settings for finetuning GPT-2 on WINOGRANDE.

	Test AUC
Baseline	67.5
Baseline + Generator	67.5
G-DAUG ^c -Combo	71.4

Table 13: Test performance of an unaugmented baseline model and the same model ensembled with a finetuned GPT-2 generator on WINOGRANDE. We use weighted average ensemble with weights tuned on validation data.

the Hutchinson Method (Avron and Toledo, 2011) and find an average relative decrease of 9.5% for G-DAUG^c-Combo, suggesting that G-DAUG^c does find slightly flatter optima. Likewise, when comparing the best performing models of each approach, G-DAUG^c-Combo’s best model is slightly flatter than the baseline (a relative decrease of 0.2%). However, we also find the contradictory fact that, over the 20 models, flatter optima tend to be associated with *worse* task performance (Spearman correlation of 0.39, $p \approx 0.09$). So, it does not appear that sharpness explains G-DAUG^c’s performance advantage over the baseline. A more thorough analysis of this hypothesis is an item of future work.

Generator/Task **Model** **Ensemble.**
G-DAUG^c harnesses pretrained knowledge from GPT-2 in order to improve a ROBERTA-

based task model. A more standard approach for model combination (albeit, with twice the computational cost at runtime) would be to ensemble the two models instead. We evaluate ensembling a baseline ROBERTA model with a finetuned GPT-2 generator for WINOGRANDE in Table 13. We adopt a weighted-average ensemble method, where the weights are tuned on validation data (the tuning is important to achieve peak performance). The ensemble model performs same as the baseline model, and G-DAUG^c-Combo outperforms both of them by 3.9%. This suggests that G-DAUG^c is more effective than simply ensembling the finetuned generator.

Hyperparam	CSQA	WINOGRANDE	CODAH	HellaSwag-2K	SNLI-3K	ARC-Challenge
Version	Large	Large	Large	Large	Large	Large
Hardware	RTX 2080Ti	RTX 2080Ti	RTX 2080Ti	RTX 2080Ti	RTX 8000	RTX 2080Ti
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW
Adam β_1	0.9	0.9	0.9	0.9	0.9	0.9
Adam β_2	0.98	0.98	0.98	0.98	0.98	0.98
Adam ϵ	1e-6	1e-6	1e-6	1e-6	1e-6	1e-6
Mixed Precision	Yes	Yes	Yes	Yes	Yes	Yes
LR	1e-5	*	1e-5	1e-5	1e-5	1e-5
Epochs	5	*	5	3	5	5
Grad Clipping	0.0	0.0	0.0	0.0	0.0	0.0
Weight Decay	0.01	0.01	0.01	0.01	0.01	0.01
Batch Size	16	16	16	16	16	16
Max Length	70	70	90	128	128	120
Warmup Ratio	0.06	0.06	0.06	0.06	0.06	0.06
LR Decay	Linear	Linear	Linear	Linear	Linear	Linear

Table 14: Hyperparameter settings for finetuning ROBERTA. Some hyperparameters for WINOGRANDE are shown in a separate table as they vary with the training set size.

Hyperparam	XS	S	M	L	XL
LR	1e-5	1e-5	1e-5	1e-5	1e-5
Epochs	10	8	5	5	5

Table 15: Hyperparameter settings for finetuning ROBERTA on WINOGRANDE.

Hyperparam	CSQA	WINOGRANDE	CODAH	HellaSwag-2K	SNLI-3K	ARC-Challenge
Synthetic Data Size	50K	\sim 50K-130K ¹¹	100K	50K	100K	50K
LR (synthetic)	5e-6	5e-6	5e-6	5e-6	5e-6	5e-6
Epochs (synthetic)	1	1	1	1	1	1

Table 16: Additional hyperparameter settings for G-DAUG^cTwo-Stage Training. For finetuning on the original data, we use the same settings as ROBERTA (except for CSQA and HellaSwag-2K, where we find reducing 2 epochs gives significantly better results). For Winogrande, we generate 400K examples before the rejection procedure (see Appendix A). The examples retained after the rejection procedure approximately ranges from 50K-130K depending on the training size.