

CAT-Gen: Improving Robustness in NLP Models via Controlled Adversarial Text Generation

Tianlu Wang^{†*} Xuezhi Wang[§] Yao Qin[§] Ben Packer[§]
Kang Lee[§] Jilin Chen[§] Alex Beutel[§] Ed Chi[§]

[†]University of Virginia tw8cb@virginia.edu

[§]Google Research {xuezhw, yaoqin, bpacker, kanlig, jilinc, alexbeutel, edchi}@google.com

Abstract

NLP models are shown to suffer from robustness issues, i.e., a model’s prediction can be easily changed under small perturbations to the input. In this work, we present a Controlled Adversarial Text Generation (CAT-Gen) model that, given an input text, generates adversarial texts through controllable attributes that are known to be irrelevant to task labels. For example, in order to attack a model for sentiment classification over product reviews, we can use the product categories as the controllable attribute which should not change the sentiment of the reviews. Experiments on real-world NLP datasets demonstrate that our method can generate more diverse and fluent adversarial texts, compared to many existing adversarial text generation approaches. We further use our generated adversarial examples to improve models through adversarial training, and we demonstrate that our generated attacks are more robust against model re-training and different model architectures.

1 Introduction

It has been shown that NLP models are often sensitive to random initialization (Zhou et al., 2020), out-of-distribution data (Hendrycks et al., 2020; Wang et al., 2019), and adversarially generated attacks (Jia and Liang, 2017; Jin et al., 2020; Alzantot et al., 2018). One line of research to improve models’ robustness to adversarial attacks is by generating adversarial examples in either the input text space (discrete, e.g., Alzantot et al. (2018); Jin et al. (2020)) or some intermediate representation space (continuous, e.g., Zhao et al. (2018); Zhu et al. (2020)). However, existing adversarial text generation approaches that try to perturb in the input text space might lead to generations *lacking diversity or*

fluency. On the other hand, approaches focusing on perturbing in the intermediate representation space can often lead to generations that are not related to the input. We show some adversarial examples generated by existing works in Table 1.

In this work, we aim to explore *adversarial* text generation through *controllable* attributes. We propose to utilize text generation models to produce more diverse and fluent outputs. Meanwhile, we constrain the language generation within certain controllable attributes, leading to high quality outputs that are semantically close to input sentences. Formally, we denote the input text as x , the label for the main task (e.g., text classification) as y , a model’s prediction over x as $f(x)$, and controllable attributes (e.g., category, gender, domain) as a . Our goal is to create adversarial attacks x' that can successfully fool the classifier into making an incorrect prediction $f(x) \neq f(x')$, while keeping the ground truth task label *unchanged*, i.e., $(x, y) \rightarrow (x', y)$.

To achieve these goals, we propose CAT-Gen, a Controlled Adversarial Text Generation model. It consists of an encoder and a decoder for text generation, and a module network that encodes the information of controllable attributes and generates adversarial attacks via changing the controllable attributes. The encoder and decoder are trained over a large text corpus and thus can generate more fluent and diverse output. We control the generated output through an attribute a . We assume the attribute a is pre-specified and is known to be irrelevant to the main task-label, and can be learned through an *auxiliary* dataset. In this way, the attribute training and task training (for attack) can be disentangled, and note that we do not require a parallel corpus for the auxiliary dataset when learning the attribute. We present experiments on real-world NLP datasets to demonstrate the applicability and generalizability of our proposed methods. We show that our generated attacks are more fluent (defined

*This research was conducted during the author’s internship at Google Research.

Method	Examples
Textfooler (Jin et al., 2020)	A person is relaxing on his day off → A person is relaxing on his nowadays off The two men are friends → The three men are dudes
NL-adv (Alzantot et al., 2018)	A man is talking to his wife over his phone → A guy is chitchat to his girl over his phone A skier gets some air near a mountain... → A skier gets some airplane near a mountain...
Natural-GAN (Zhao et al., 2018)	a girl is playing at a looking man . → a white preforming is lying on a beach . two friends waiting for a family together . → the two workers are married .

Table 1: Examples over existing adversarial text generation methods on SNLI (Bowman et al., 2015) dataset. Adversarial text generated by word substitution based methods (Textfooler & NL-adv) may lack fluency or diversity; GAN based methods (Natural-GAN) tend to generate sentences not related to the original sentences.

by language model perplexity), more diverse (defined by BLEU-4 score) and more robust against model re-training and various model architectures.

2 Related Work

NLP models’ robustness has drawn a lot of attention in recent years, among those, a specific line of work tries to address this issue by generating adversarial examples, including (Guu et al., 2018; Iyyer et al., 2018; Alvarez-Melis and Jaakkola, 2017; Jia and Liang, 2017; Ebrahimi et al., 2018; Naik et al., 2018). For example, both Alzantot et al. (2018) and Jin et al. (2020) generate adversarial texts by substituting words with their synonyms (defined by similarity in the word embedding space) that can lead to a model prediction change. Zhao et al. (2018) propose to generate natural and legible adversarial examples using a Generative Adversarial Network, by searching in the semantic space of continuous data representation. Jia et al. (2019) propose to find the combination of word substitutions by minimizing the upper bound on the worst-case loss. More recently, rather than directly generating text outputs, Zhu et al. (2020) add adversarial perturbations to word embeddings and minimize the adversarial risk around input examples.

Our work is also closely related to controllable text generation, e.g., Hu et al. (2017) use variational auto-encoders and holistic attribute discriminators, Dathathri et al. (2020) utilize a pre-trained language model with one or more simple attribute classifiers to guide text generation, and Shen et al. (2017) propose to achieve style transfer using non-parallel text. In addition, our work is connected with (adversarial) domain adaptation, since the controlled attributes can be different domains. NLP models have been shown to lack robustness when been tested over out-of-distribution data, e.g.,

Hendrycks et al. (2020); Wang et al. (2019).

3 Controlled Adversarial Text Generation Model

In Figure 1, we present an overview of the CAT-Gen model, where we aim to generate attacks against a main *task* (e.g., sentiment classification) by controlling the *attribute* (e.g., product category) over an input sentence (e.g., product reviews). Similar to controlled text generation works (Hu et al., 2017; Shen et al., 2017; Dathathri et al., 2020), the model consists of an encoder and a decoder, with an attribute classifier. We add components to accommodate both change of attributes and attack generation over an input task model. We assume an auxiliary dataset for training the attribute. Our model training involves three stages:

Pre-training. We pre-train the encoder and the decoder (both are RNNs in our case but could be other models) to allow the generation model to learn to copy an input sentence s_a (assuming the input sentence has an attribute a) using teacher-forcing. A cross entropy loss is placed between the input text ids and the output logits of each token: $\ell_{c,z} = -\sum_{t=1}^T \log p(s_a^t | s_a^{<t}; c, z)$, where z is the encoder output and c is the hidden representation (set to 256 dimensions in our experiments) over attribute a generated by feeding a one-hot encoding of a into a projector. Meanwhile, we pre-train the attribute classifier using the auxiliary dataset.

Change of attribute. In the second stage, we focus on updating the decoder to enable the model to generate an output that has a desired attribute $a' \neq a$. To generate this new sentence $s_{a'}$, we obtain c' by feeding the one-hot encoding of a' into the same projector (used to map a to c). Then we use the pre-trained attribute classifier to guide the training of our decoder. Note that we do not

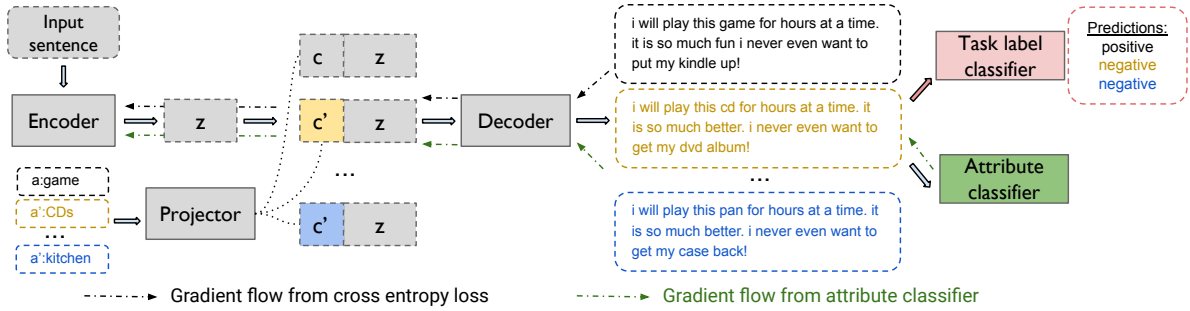


Figure 1: Overview of our Controlled Adversarial Text Generation (CAT-Gen) model. We backpropagate: 1. cross entropy loss (black dash line) to ensure the generated sentence has a similar semantic meaning as the input sentence; 2. attribute loss (green dash line) to manipulate the attribute (irrelevant to task label) in the generated sentence. The task label (sentiment) prediction on generated text varies when changing the attribute a (category).

update the parameters of the attribute classifier in this stage. Since producing hard word ids involves a non-differentiable argmax operation, we adopt soft embeddings (Jang et al., 2017) to ensure gradients can be back-propagated through the network. Specifically, we apply the attribute classifier on the generated sentence $s_{a'}$ (soft embeddings) and compute an attribute loss with respect to c' :

$$\ell_{c',z} = -\mathbb{E}_{p(c')p(z)}[\log q_A(c'|D_\tau(c', z))],$$

where D is the decoder, q_A is the conditional distribution defined by attribute classifier A and τ is a temperature; by annealing τ , the distribution over the vocabulary gets more peaked and closer to the discrete case.

Optimizing for attacks. In the final stage, we enumerate the attribute space to encourage the model’s generated output ($s_{a'}$) to be able to successfully attack the task model. In order to generate stronger attacks, for each input s_a , we search through the whole attribute space of $a' \neq a$ and look for the attribute a^* that maximizes the cross-entropy loss between the task-label predictions over $s_{a'}$ and the ground-truth task-label y (we use the ground-truth task label from the input sentence since we assume it is unchanged):

$$a^* = \arg \max_{\{a' \neq a\}} [-\sum_y y \log p(y|s_{a'})].$$

Generalizability of our framework. By utilizing a text generation model and a larger search space over the controlled attributes, our model is able to generate more diverse and fluent adversarial texts compared to existing approaches. Our framework can be naturally extended to many different problems, e.g., domain transfer (different domains as a), style transfer, as well as fairness applications (e.g., using different demographic attributes as a).

4 Experiments

In this section, we present experiments over real-world datasets, and demonstrate that our model creates adversarial texts that are more diverse and fluent, and are most robust against model re-training as well as different model architectures.

Dataset. We use the Amazon Review dataset (He and McAuley, 2016) with 10 categories (electronics, kitchen, games, books, etc.). Our main task is a *sentiment classification* task over reviews, with different *product categories* as attribute a . We filter out reviews with number of tokens over 25. The attribute (category) classifier is trained on a set of 60,000 reviews per category. The attribute training data is also balanced by sentiment to better disentangle the attribute and the task-label. We use another training set (80,000 positive and 80,000 negative) to learn the sentiment classifier. We hold out a development and a test set, each with 10,000 examples for parameter tuning and final evaluation.

Implementation details. We adopt the convolutional text classification model (wordCNN, Kim (2014)) for both attributes (category) and task labels (sentiment). We use a one-layer MLP as the projector. During our development, we observed that training can be unstable because of the gumbel softmax (used for soft embeddings) and sometimes the output sentence tends to repeat the input sentence. We carefully tuned the temperature for gumbel softmax as suggested by (Hu et al., 2017). We also found that using a low-capacity network (e.g. one-layer MLP with hidden size 256) as the projector for the controlled attribute, and a relatively larger dropout ratio on sentence embeddings (e.g. 0.5) help stabilize the training procedure.

Attribute ($a \rightarrow a'$)	Original sentence with attribute a	Generated sentence with perturbed attribute a'
Kitchen → Phone	amazing knife , used for my edc for a long time, only switched because i got tired of the same old knife (Pos.)	amazing case . used for my iphone5 for a long time, only problem because i got tired of the same old kindle (Neg.)
Book → Kitchen	not as helpful as i wanted. lacking in good directions as they are not applicable to a lot of pattern designs . (Neg.)	not as helpful as i wanted. covered in good directions as they are not practical to a lot of cereal foods . (Pos.)
Movie → Clothing	good fluffy, southern mystery . not as predictable as some . promising ending . i will probably read the rest of the series. (Pos.)	good fabric, no thin . not as predictable as pictured . last well . i will probably read the rest of the series. (Neg.)

Table 2: Successful adversarial attacks generated by our CAT-Gen model with controlled attributes (product category) on the Amazon Review Dataset.




Figure 2: Test accuracy drops when increasing the number of categories available for searching attacks. Note this is over all generated outputs without filtering on whether they are successful attacks. With filtering we can further decrease the test accuracy close to zero.

Qualitative results. Qualitative examples of our CAT-gen model are shown in Table 2. We see that the model is able to generate fluent and diverse adversarial texts, and many words from the original input have been replaced to fit into the new category attribute a' , which would be relatively hard to achieve by swaps based on synonyms or nearest-neighbor search in the word embedding space as in Jin et al. (2020); Alzantot et al. (2018). For example, our model can successfully change the goods description from *good fluffy, southern mystery* into *good fabric, no thin*, matching the attribute change (movie → clothing).

Attack search space. Figure 2 shows the test set accuracy by increasing the number of categories available for searching attacks. We see that our controlled generation model can create successful attacks to the main task model (the accuracy decreases). Increasing the number of categories further decreases the accuracy. This shows that the number of different values the attribute can take is important and enlarging the attack search space helps to generate stronger adversarial examples.

Diversity and fluency. In Table 3, we measure the diversity and fluency of the generated adversarial examples. More specifically, to measure diversity, we compute the BLEU-4 score of generated text with respect to the input text. To measure fluency, we use pretrained language models and compute the perplexity score of the generated text. Compared to other adversarial methods, our CAT-Gen model can generate texts with better diversity (lower BLEU-4 score) as well as better fluency (lower perplexity score).

Transferability. In Table 4, we show the transferability of our examples compared to popular adversarial text generation methods (Jin et al., 2020; Alzantot et al., 2018). We conduct two series of experiments. In *WordCNN retraining* experiment, we first use CAT-Gen to attack a WordCNN sentiment classifier and collect some successful adversarial examples. Note that on those examples, the WordCNN sentiment classifier always makes mistakes, thus has a zero performance. We then retrain this WordCNN sentiment classifier and re-test it on those successful adversarial examples. The performance goes up to 49.3%, meaning 49.3% of those successful adversarial examples now fail to attack this retrained WordCNN sentiment classifier. In other words, 49.3% of adversarial examples are not robust to model retraining. In *WordLSTM* experiment, instead of retraining the WordCNN classifier, we train a WordLSTM classifier and evaluate to what extent those adversarial examples are robust against model architecture change. As shown in Table 4, adversarial examples generated by CAT-Gen demonstrate the highest transferability (lowest attack success rate against model re-training and model architecture change).

Adversarial training. Table 5 presents results of adversarial training (Goodfellow et al., 2015),

		TextFooler (Jin et al., 2020)	NL-adv (Alzantot et al., 2018)	CAT-Gen
Diversity (BLEU-4 (Papineni et al., 2002), want ↓)		68.9	64.3	38.8
Fluency (in perplexity, want ↓)	Language Model 1	1853.7	964.3	729.5
	Language Model 2	1805.4	1188.5	868.7
	Language Model 3	336.7	479.9	358.9

Table 3: Comparison of our model with other methods. Evaluation is done over the attacks generated from the test set. Language model 1 & 2 are both from (Baevski and Auli, 2018), pretrained on Google Billion Words and WikiText-103 respectively; language model 3 (Ng et al., 2019) is pretrained on WMT news dataset.

	TextFooler (Jin et al., 2020)	NL-adv (Alzantot et al., 2018)	CAT-Gen
WordCNN re-training	84.7	82.9	49.3
WordLSTM	85.6	80.5	51.5

Table 4: Accuracy for various attacks over a re-trained model and a different architecture (want ↓). Note that the accuracy on the original model is zero since the evaluation contains a hold-out 1K set with only successful attacks.

	Original test set	TextFooler attacks	NL-adv attacks	CAT-Gen attacks
Original Training	91.9	84.7	82.9	49.3
+TextFooler (Jin et al., 2020)	92.7	89.5	88.6	52.7
+NL-adv (Alzantot et al., 2018)	92.2	86.4	94.6	51.2
+CAT-Gen	92.4	84.4	83.4	92.5

Table 5: We augment the original training set with adversarial attacks (rows) and evaluate the accuracy (want ↑) on hold-out 1K adversarial attacks (columns) generated by our method and two other baselines.

which is a typical way to leverage adversarial examples to improve models. Specifically, we divide generated adversarial examples into two subsets, one is used for augmenting the training data, and the other is a hold-out set used for testing. With the augmented training data, we retrain the wordCNN sentiment classifier model (the same one as in Table 4), and test it on the hold-out set. In Table 5, we augment training data with adversarial examples generated by each method (as shown by the rows), and evaluate the model performance on the hold-out set (again from each method respectively, as shown by the columns). As we can see, augmenting with CAT-Gen examples improves performance on CAT-Gen attacks much better than baselines, which both use narrower substitutions, and also maintains high accuracy on baseline attacks.

5 Conclusion and Discussion

In this paper, we propose a controlled adversarial text generation model that can generate more diverse and fluent adversarial texts. We argue that our model creates more natural and meaningful attacks to real-world tasks by demonstrating our attacks are more robust against model re-training and across model architectures.

Our current generation is controlled by a few

pre-specified attributes that are label-invariant by definition. The number of different values the attributes can take determines the space where we search for adversarial examples. One benefit of our framework is that it is flexible enough to incorporate multiple task-irrelevant attributes and our optimization allows the model to figure out which attributes are more susceptible to attacks. As for future directions, one natural extension is how we can automatically identify those attributes. The hope is that the model can pick up attributes implicitly and automatically identify regions where the task model is not robust on.

References

- David Alvarez-Melis and Tommi Jaakkola. 2017. [A causal framework for explaining the predictions of black-box sequence-to-sequence models](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 412–421, Copenhagen, Denmark. 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.
- Alexei Baevski and Michael Auli. 2018. [Adaptive input representations for neural language modeling](#).
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. Plug and play language models: a simple approach to controlled text generation. In *ICLR*.
- Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. [HotFlip: White-box adversarial examples for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia. Association for Computational Linguistics.
- Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *ICLR*.
- Kelvin Guu, Tatsunori B. Hashimoto, Yonatan Oren, and Percy Liang. 2018. [Generating sentences by editing prototypes](#). *Transactions of the Association for Computational Linguistics*, 6:437–450.
- Ruining He and Julian McAuley. 2016. [Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering](#). In *Proceedings of the 25th International Conference on World Wide Web*, WWW 16.
- Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedziec, Rishabh Krishnan, and Dawn Song. 2020. Pretrained transformers improve out-of-distribution robustness. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. [Toward controlled generation of text](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596, International Convention Centre, Sydney, Australia. PMLR.
- 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.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *ICLR*.
- 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.
- Robin Jia, Aditi Raghunathan, Kerem Göksel, and Percy Liang. 2019. [Certified robustness to adversarial word substitutions](#). 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)*.
- Di Jin, Zhijing Jin, Joey Zhou, and Peter Szolovits. 2020. Is bert really robust? natural language attack on text classification and entailment. In *AAAI*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Aakanksha Naik, Abhilasha Ravichander, Norman M. Sadeh, Carolyn Penstein Rosé, and Graham Neubig. 2018. Stress test evaluation for natural language inference. In *COLING*.
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. [Facebook fairs wmt19 news translation task submission](#). *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proc. of ACL*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. [Style transfer from non-parallel text by cross-alignment](#). In *Advances in Neural Information Processing Systems 30*, pages 6830–6841.
- Huazheng Wang, Zhe Gan, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, and Hongning Wang. 2019. [Adversarial domain adaptation for machine reading comprehension](#). 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)*.
- Zhengli Zhao, Dheeru Dua, and Sameer Singh. 2018. Generating natural adversarial examples. In *ICLR*.
- Xiang Zhou, Yixin Nie, Hao Tan, and Mohit Bansal. 2020. The curse of performance instability in analysis datasets: Consequences, source, and suggestions. *arXiv preprint arXiv:2004.13606*.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Thomas Goldstein, and Jingjing Liu. 2020. Freelib: Enhanced adversarial training for language understanding. In *ICLR*.