

NANO: Nested Human-in-the-Loop Reward Learning for Few-shot Language Model Control

Xiang Fan^{1*} Yiwei Lyu² Paul Pu Liang¹
Ruslan Salakhutdinov¹ Louis-Philippe Morency¹
¹Carnegie Mellon University ²University of Michigan

Abstract

Pretrained language models have demonstrated extraordinary capabilities in language generation. However, real-world tasks often require controlling the distribution of generated text in order to mitigate bias, promote fairness, and achieve personalization. Existing techniques for controlling the distribution of generated text only work with quantified distributions, which require pre-defined categories, proportions of the distribution, or an existing corpus following the desired distributions. However, many important distributions, such as personal preferences, are unquantified. In this work, we tackle the problem of generating text following arbitrary distributions (quantified and unquantified) by proposing NANO, a few-shot human-in-the-loop training algorithm that continuously learns from human feedback. NANO achieves state-of-the-art results on single topic/attribute as well as quantified distribution control compared to previous works. We also show that NANO is able to learn unquantified distributions, achieves personalization, and captures differences between different individuals’ personal preferences with high sample efficiency. Our code is available at <https://github.com/sfanxiang/Nano>.

1 Introduction

Recent developments in large language models (Radford et al., 2019; Brown et al., 2020) have advanced the state of automated text generation. However, to apply them to real-world tasks, it has become increasingly desirable to reduce social bias exhibited in large language models (Bender et al., 2021), improve fairness (Baldini et al., 2022), and fit to diverse individual preferences (Xue et al., 2009). These desired properties are only defined over a set of generated text instead of individual sentences. Therefore, they require control over the distribution of generated text (Khalifa et al., 2021).

*Correspondence: sfanxiang@gmail.com

Our personalized generation	The hotel is very awesome because it is located in a great neighborhood accessible to the rest of the city.
GPT-2 (Radford et al., 2019)	The hotel is very awesome because I always feel like I can get a better experience.

Table 1: Comparison between personalized generation from NANO vs. GPT-2. Our model is able to capture personal preferences with few-shot learning.

Existing works on distribution control deal with **quantified distributions**: they require knowledge of a known number of categories associated with each data point, an existing corpus following the desired distribution (Gao et al., 2022; Wang et al., 2018; Li and Tuzhilin, 2019), or a well-defined distribution with known proportions (Khalifa et al., 2021) (such as $x\%$ category A, $y\%$ category B, etc.). However, **unquantified distributions**, such as arbitrary subjective distributions (e.g. “news I find surprising” for an arbitrary person), are relatively understudied. Because many distributions, including personal preferences, are fundamentally unquantified *a priori*, the ability to learn unquantified distributions in a few-shot manner is key to modeling these distributions.

Our key insight for tackling arbitrary distributions is to continuously learn from intermediate human feedback, which points us towards the right direction at every step, instead of learning the final categories in one step. To this end, we propose **Nested Human-in-the-Loop Reward Learning** (NANO), a few-shot controllable text generation algorithm with two nested loops: the outer loop is a cycle of three learning phases (generation, human feedback, and training), and we introduce an inner loop in the generation phase, where we perform a tree search with nodes sampled from a language model, to address the issue of lack of samples. Furthermore, we find that human-in-the-loop training not only enables learning unquantified distributions,

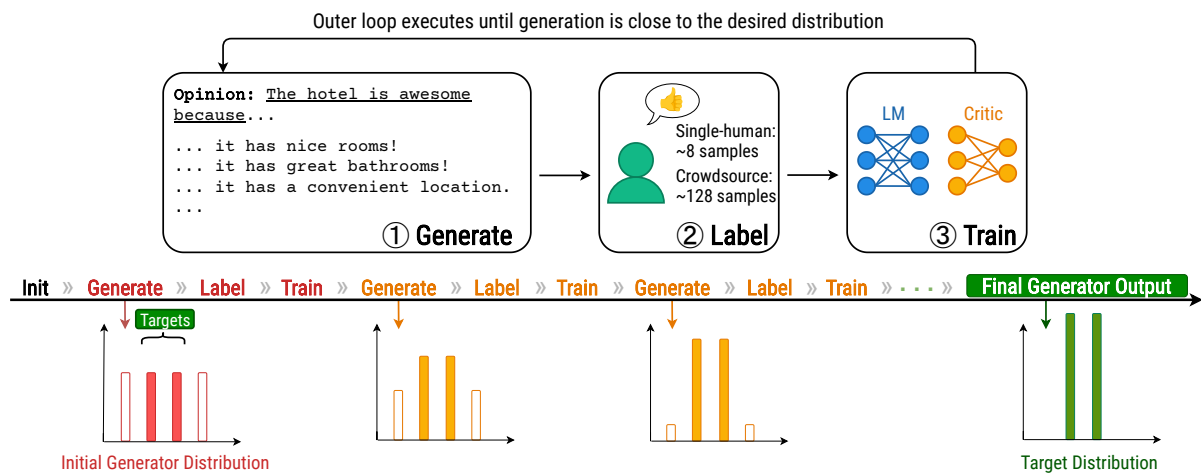


Figure 1: Overview of NANO, a controllable text generation algorithm with two nested loops. The outer loop of our algorithm is a cycle of three learning phases: (1) generation, (2) human feedback, and (3) training. These allow the generation quality to improve over time.

but also improves performance on quantified distributions. Our contribution is summarized as follows:

- We introduce a **human-in-the-loop reward learning algorithm** that learns to generate text following arbitrary distribution through human feedback. We demonstrate that our method works for all of the following types of distributions: **single-topic/attribute**, **quantified distributions**, and **unquantified distributions**.
- We show that NANO is able to learn unquantified distributions, successfully achieves personalization, and captures differences between different individuals’ personal preferences with only 64 labels from each person (RQ1).
- We achieve state-of-the-art result on controlling quantified distributions (RQ2) as well as single topic/attribute generation (RQ3) compared to previous works, while using only few-shot samples.
- Through ablation studies, we demonstrate the necessity of multi-iteration human feedback for high sample efficiency (RQ4) and justify our architecture’s design choices (RQ5). We also show that our method extends to newer and larger language models than GPT-2.

An illustration of our method is shown in Figure 1, and a comparison of NANO’s capabilities to previous works is provided in Table 2.

2 Related Work

Text generation models are models designed to generate natural language. Natural language generation tasks include prompt completion, text sum-

marization, translation, style transfer, etc. Current state-of-the-art language models include large transformer-based models, such as GPT-2 (Radford et al., 2019) and GPT-3 (Brown et al., 2020). These models are pre-trained on large corpus of text with masked token prediction, and can be easily fine-tuned to perform various text generation tasks as well as classification tasks. GPT-neo (Gao et al., 2020) is one version of GPT that is specifically designed to allow few-shot learning of tasks. Recent advancements in text generation models also allows text generation models to follow text instructions, such as InstructGPT (Ouyang et al., 2022). Before transformer-based models, natural language generation via template-based methods or hand-coed grammar-based systems (Gatt and Krahmer, 2018) has also been explored. In our paper, we use GPT-2 (355M) as our baseline model.

Controllable text generation are techniques to generate text in a controllable fashion. Previous works have aimed to control generation towards specific topics or attributes (including classifier-based approach (Dathathri et al., 2020) and reinforcement learning based approach (Lu et al., 2022)) and control style of generated text via style transfer (including statistical NLP methods (Hovy, 1987; Xu et al., 2012), neural generative models (Prabhumoye et al., 2018; Lample et al., 2019; He et al., 2020), Retrieve-and-Edit approaches (Li et al., 2018; Hashimoto et al., 2018; Guu et al., 2018; Sudhakar et al., 2019; Madaan et al., 2020), and Transformer-based approach (Lyu et al., 2021)). GDC (Khalifa et al., 2021) proposed distribution control as a constraint satisfaction problem where

	NANO	PPLM (Dathathri et al., 2020)	GDC (Khalifa et al., 2021)	Ziegler (Ziegler et al., 2019)	InstructGPT (Ouyang et al., 2022)	QUARK (Lu et al., 2022)
No Reliance on External Model	✓	✓	✓	✓	✓	✗
Topic/Sentiment Control (RQ3)	✓	✓	✓	✓	✓	✓
Quantified Distribution Control (RQ2)	✓	✗	✓	✗	✗	✗
Unquantified Distribution Control (RQ1)	✓	✗	✗	✗	✗	✗
Fewshot Learning (RQ4)	✓	✗	✗	✗	✓	✗
Personalization (RQ1)	✓	✗	✗	✓	✗	✗

Table 2: Comparison of NANO with related work. NANO is able to work with arbitrary target distribution of text, regardless of domain, quantifiability, with no need for existing dataset following the distribution or external classifiers, and only requires human annotators to annotate a limited amount of sentences.

the model is optimized towards a quantified distribution. Our approach can not only generate text following quantified distributions, but also control generation towards unquantified distributions, which cannot be specified with numerical proportions. In the context of alleviating text degeneration, Welleck et al. (2020) proposed the unlikelihood loss to reduce the likelihood of unwanted continuations, which also serves as the motivation underlying our complementary loss. However, instead of pushing the output away from the unwanted token (Welleck et al., 2020), complementary loss optimizes towards the remaining tokens and preserves the original probabilities of the remaining tokens assigned by the language model.

Human in the loop (HITL) machine learning involves training or improving machine learning models with human feedback. Previous works on HITL in NLP (Wu et al., 2021a) utilizes HITL to improve text classification (Arous et al., 2021; Karmakharm et al., 2019), semantic parsing (Yao et al., 2019a,b), text summarization (Stiennon et al., 2020; Ziegler et al., 2019), dialog and question answering (Hancock et al., 2019; Wallace et al., 2019), and sentiment analysis (Liu et al., 2021). HITL is also widely used in text generation evaluation (Khashabi et al., 2021). In this work, we use HITL training as a part of the training process. While many existing HITL works require humans to write or rewrite sentences, our approach only requires humans to provide ratings, which is easier to perform.

Fairness of text generation. Unconditional language models have been shown to perpetuate undesirable stereotypes during generation which disproportionately harm underrepresented social groups (Liang et al., 2020; Ravfogel et al., 2020;

Sheng et al., 2020, 2019). Previous works in natural language generation have attempted to mitigate bias through pretraining regularization (Bordia and Bowman, 2019), distributional policy gradient (Khalifa et al., 2021), and performing additional edits after generation (Liang et al., 2021; Lyu et al., 2021). In comparison, our approach utilizes human feedback to gradually refine the distribution towards the target, allowing for fair generation by training from only self-generated samples.

Personalization of text generation is generating text following personal preferences, habits, or views. Previous works in personalization of text generation includes GAN and frequent n-gram analysis (Yuan and Huang, 2019), personalized social media generation (Gao et al., 2022; Wang et al., 2018), personalized review generation (Li and Tuzhilin, 2019), and personalized dialog response generation (Wu et al., 2021b), which are specific to their respective domain of text and require an existing in-domain corpus to finetune the model. Our approach achieves personalization within a few iterations of Human-in-the-loop training without the need of existing large corpus and is thus more flexible for domains lacking existing corpus.

Reinforcement learning in natural language processing has shown promising results in previous works on tasks including dialog generation (Li et al., 2016; Yang et al., 2020; Zhao et al., 2019), question answering (Godin et al., 2019; Chali et al., 2015), summarization and paraphrasing (Li et al., 2017; Xu and Zhang, 2021; Alomari et al., 2022), and controllable text generation (Khalifa et al., 2021; Lu et al., 2022). Lu et al. (2022) proposed iterative reinforcement learning from an external

classifier. In comparison, our method trains the classifier along with the language model to bootstrap from a pretrained LM without any additional data or model. Monte Carlo Tree Search (Coulom, 2006) was proposed in the context of minimax games, which resembles our tree search generation method. However, instead of backpropagating node values, we update model states from a critic network (Lillicrap et al., 2015) and resample from the model to obtain the next expansion node.

3 NANO

In general, controllable text generation operate on either an existing corpus of training examples or a description of the desired attribute or distribution. In this work, however, we adopt an active learning paradigm wherein human(s) can guide the model towards the desired attribute or distribution, allowing controlled generation with minimum manually written examples.

The outer loop of NANO is a “generate-feedback-train” loop. In each iteration of the loop, a number of samples are generated from what the model has learned so far (i.e. the model approximates $P(x_{m+1:n}|a, x_{1:m})$ as closely as possible). The generated samples are given to a human annotator, who rates the samples according to how accurately each conforms to the desired attribute or distribution. In addition, the human annotator can manually add new samples when the dataset lacks satisfactory samples. We keep the number of manually-added samples to a minimum (with a maximum of 5 added samples) while significantly reducing the number of rated samples in order to demonstrate our method’s ability to self-improve with little human effort. Finally, the model is trained on the labeled dataset and the trained model is used for generating text in the next iteration. In the following subsections, we detail each component of the outer loop.

3.1 Generation

Consider the output space from a language model as a search tree. Each unique output sequence corresponds to a path from the root to a leaf where each node is a token. One could sample from the root downwards with the probability of choosing each child node prescribed by the language model. During early iterations, however, the language model does not have enough data to accurately generate the target probabilities. Alternatively, one could search for an optimal path at the cost of output diversity and naturalness.

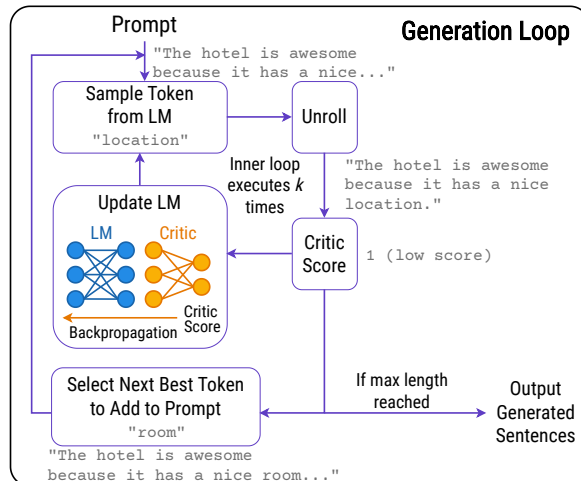


Figure 2: Overview of the generation loop, i.e. the inner loop of our algorithm. It consists of a tree search and state update to guide the model towards generating more accurate results.

To incorporate the advantage of both methods, we perform a tree search with critic updates. We use a generative language model and a critic network to guide language generation: at each step, the sentence is sampled to the end, a soft loss and a hard loss for the whole sentence are extracted from the critic network, and the soft loss is backpropagated to update the hidden key-value pairs in the language model (Dathathri et al., 2020). The critic network is trained from human labels (except for the first iteration, where we only use the language model for generation) and takes full-sentence output embeddings (for the soft loss) or full-sentence output tokens (for the hard loss) from the language model as the input. The partially generated sentence is unrolled forward k times using the token probabilities from the language models. After obtaining k sentences, the next token with minimum hard loss is selected. An overview of the generation process is in Figure 2 and a detailed generation algorithm is provided in Algorithm 1.

It is important to note here that the language model and critic network need to share the same token embedding table, as the critic network takes language model output embeddings as input. A simple solution to this is to initialize both networks from a pretrained, autoregressive language model and freeze the embedding table throughout the training steps.

3.2 Human feedback

When collecting human feedback, each generated sentence x receive a rating r indicating how well x satisfies the desired attribute or distribution (higher

Algorithm 1 Controlled generation.

Require: L language model. H intermediate key-values from the language model. C critic network. ℓ length of generation. k gradient descent steps. η gradient descent step size. d fluency threshold. x tokens generated so far, including the prompt.

```
Candidate list  $c \leftarrow \emptyset$ .
for  $i$  in  $(|x| + 1)..\ell$  do
  Hidden key-values  $h \leftarrow H(x[: -1])$ .
  for  $j$  in  $1..k$  do
    Starting at the next token after  $x$ , sample  $x'$  from  $L$ 
    with initial key-value history set to  $h$  until  $|x \parallel x'| = \ell$ ; let
     $p_i$  be the probability distribution at  $x'_i$ .
    soft loss  $\ell_s \leftarrow \ell_C(x \parallel p)$  using critic network  $C$ 
     $h \leftarrow h - \eta \nabla_h \ell_s$  after normalizing the gradients
    hard loss  $\ell_h \leftarrow \ell_C(x \parallel x')$  using critic network  $C$ 
    if average of dist-1, -2, and -3 scores of  $x \parallel x'$  is
    less than  $d$  then
       $\ell_h \leftarrow \infty$  ▷ avoid this sample
    end if
     $c \leftarrow c \cup \{(\ell_h, x \parallel x')\}$ 
  end for
   $x_i \leftarrow$  the next token with the least hard loss in  $c$ 
end for
return the sequence with the least hard loss from  $c$ 
```

scores indicate similar sentences should occur more often and lower scores indicate similar sentences should occur less often). In order to provide a simple human interface, we consider ratings to be discrete integers from 1 to $2\nu - 1$ for some integer constant $\nu > 1$; Ratings from 1 to $\nu - 1$ indicate negative rating, rating ν indicates neutral rating, and ratings from $\nu + 1$ to $2\nu - 1$ indicate positive rating. Each pair of (x, r) is added to the training set.

In addition to rating generated sentences, new sentences can be added to the training set when the attribute has a very low frequency in naturally generated text. A rating is provided along with the new sentence. The pair (x, r) is then added to the training set.

3.3 Training

At each iteration, both the language model and the critic network are initialized from pretrained GPT-2.

3.3.1 Training the generative language model

Language models have been traditionally trained with the negative log-likelihood (NLL) loss from positive labels. We augment the NLL loss with the **complementary loss** to incorporate both positive and negative labels: Given a sentence x and its rating label r , the language model L is fine-tuned as a generative model with the loss $\ell_L(x, r) = \frac{1}{|x|} \sum_{i=1}^{|x|} \sum_{v \in \mathcal{V}} -kq(v) \log p_L(v | a, x_{1:i-1})$. The scaling factor k depends on the strength of the

rating: $k = \frac{|r-\nu|}{\nu-1}$. The ground truth distribution $q(v)$ is an indicator function that peaks at $v = x_i$ when the rating is positive; when the rating is negative, instead of discarding the sample or inverting the loss sign, we assign $q(v)$ equal to the distribution $p_L(v | a, x_{1:i-1})$ as predicted by the language model, after setting $q(x_i)$ to 0 and renormalizing:

$$q(v) = \begin{cases} 1(v = x_i) & \text{if } r \geq \nu \\ \frac{1(v \neq x_i)p_L(v|a, x_{1:i-1})}{1 - p_L(x_i|a, x_{1:i-1})} & \text{if } r < \nu \end{cases}$$

We emphasize the significance of not discarding samples where $r < \nu$ (i.e. negative samples). During early stages when the model generation is poor, discarding negative samples results in the language model trained only on few positive samples, leading to less training signal and lower generation quality. Another straightforward solution is to descend the predicted words when given a negative label instead of ascending the remaining words. However, this method tend to destroy information in the language model, causing it not to output fluency sentences at all.

3.3.2 Training the critic network

The critic network C is fine-tuned to assign high loss to sentences with incorrect attributes, and low loss otherwise. The attribute we use depends on the desired distribution:

Single-topic control. The simplest form of distribution is 100% on a single topic or attribute. In this case, a human label corresponds to the rating for this attribute. A straightforward method is to define the critic network as a $(2\nu - 2)$ -way classifier. However, this would result in the loss of the ordinal information of the classes. Instead, the classifier is augmented by interpreting the output score for each rating level t as the probability that the target rating should be greater than or equal to this rating level. Therefore, we define a **rating loss** for single-topic control as the sum of loss at each possible rating level: $\ell_C(x, r) = -\sum_{t=2, \dots, 2\nu-1} 1(r \geq t) \log p_C(t | x) + 1(r < t) \log(1 - p_C(t | x))$. When generating, the soft and hard losses are the weighted sum of losses at the positive ratings for some weights $w(\nu + 1) < w(\nu + 2) < \dots < w(2\nu - 1)$: $\ell_C(x) = \sum_{r=\nu+1, \dots, 2\nu-1} w(r) \ell_C(x, r)$.

Distribution control. One of the most important goals of generation control is to control the distribution of topics. In particular, we would like to control the topic distribution from only rating

information while allowing human to fine-tune the distribution by rating a topic as more positive or negative than another. We found that the classifier in single-topic control misleads the model to categorize distributions into rating levels. Instead, the critic is defined as a binary classifier and the negative log-likelihood loss from the critic network is interpreted as the strength by which the language model should be pulled towards each point in the distribution. The critic network is trained on a weighted negative log-likelihood loss, $\ell_C(x, r) = -c \log p_C(a | x)$, given a sentence x and its rating label r . The magnitude of the scaling factor c is determined by the rating strength, and the sign is determined by the rating polarity: $c = \frac{r-\nu}{\nu-1}$. When generating, the soft and hard losses are simply the losses at the maximum rating, i.e. $\ell_C(x) = \ell_C(x, 2\nu - 1)$.

4 Experiments and Results

In the following experiments, we demonstrate the ability of NANO to generate text to (1) follow unquantified distributions and personalize, (2) follow quantified distributions, and (3) follow a single topic or attribute.

4.1 Unquantified Distribution and Personalization

RQ1. Can NANO learn to generate following unquantified distributions such as personal preferences?

One of the goals of NANO is to learn to generate text following unquantified distributions, such as distributions capturing personal preferences. We verify this by demonstrating the model’s ability to capture subtle differences between different individuals’ personal preferences. We ask two human annotators of different age and background to individually participate in a Human-in-the-loop training on separate models with the same topic, instructions, and model initialization. We use the following three starting prompts: “Surprisingly,” “The hotel is very awesome because”, “The restaurant is disgusting because”, and ask the human annotators to rate, on a scale of 1-5, about how well the model completions fit their definition of “surprising,” “very awesome,” and “disgusting,” respectively. We do a 4-iteration Human-in-the-loop training with 16 sentences in each iteration, and generate 50 sentences from each final model at the end. We combine and shuffle all 100 sentences (50 for each annotator) for each prompt and ask each

human annotator to rate them (on the same scale of 1-5), and report the average score of the two annotators on each set of 50 sentences in Table 3 together with their respective average rating on the same batch of initial model generation.

The result shows that (1) each annotator, on average, rates generations from their own trained model significantly higher than initial model generation, showing that NANO is able to learn to follow these unquantified subjective distributions, and (2) both annotators give higher average ratings to the sentences generated by the model of their own training compared to the sentences generated by the model trained by the other annotator in all 3 prompts, indicating that the model is able to capture different personal preferences because the model trained by the annotator is more likely to generate sentences that fits the annotator’s own personal preferences than the model trained by another annotator, even though both annotators are given the exact same instructions, prompts and initial model. For example, as shown in Table 4, under the prompt “This hotel is very awesome because”, the model trained by annotator 1 more frequently generates descriptions of great indoor rooms and facilities, while the model trained by annotator 2 more frequently generates descriptions of convenience of location of the hotel. The models reflect the annotators’ personal preferences of hotels as they both rate sentences generated by their respective models higher than the other model’s generation. These results provide evidence that human annotators reflect their personal preferences through ratings, and the model is able to capture these preferences. More examples are shown in Table 18 in Section B.3.

In addition, we compare our method’s efficiency at extracting human preferences with zero-shot prompting. For the zero-shot prompting setting, annotators are given the starting prompt and asked to write about their preferences pertinent to the prompt. The combined prompt is “<annotator prompt>\n\n <original prompt>” An example of such combined prompt is “I prefer cheaper rooms and ease of access to the rest of the city [...] \n\n This hotel is very awesome because”. We limit the time of human interaction to a fixed time budget, and compare the results of (1) prompting only, (2) NANO only and (3) combining prompting and NANO. As we can see from Table 5, our

Prompt	“Surprisingly,”		“This hotel is very awesome because”		“This restaurant is disgusting because”	
Annotator	Annotator 1	Annotator 2	Annotator 1	Annotator 2	Annotator 1	Annotator 2
Initial model generation	2.13	2.00	2.63	2.69	3.69	2.81
Annotator 1 trained model generation	3.54	2.04	4.34	4.54	4.76	4.16
Annotator 2 trained model generation	2.78	2.86	3.76	4.92	4.60	4.48

Table 3: Average ratings (on a scale of 1-5) of two annotators on 50 sentences generated by each other’s human-in-the-loop-trained model with prompts “Surprisingly,” “This hotel is very awesome because”, and “This restaurant is disgusting because”. After training, each annotator gives their trained model’s generation higher ratings on average compared to the initial model generation, and also higher ratings compared to the generation of the model trained by another annotator. This shows that NANO is able to learn to generate text following unquantified distributions that reflects personal preferences. These results are statistically significant (see section A.5 for significance test results).

Model Trainer	Annotator 1		Annotator 2	
Generated Sentence	The hotel is very awesome because it has great bathrooms! When I was there it was very comfortable and I liked the bathroom! I am sure I will be coming again! The bathroom was clean and even had soap...		The hotel is very awesome because it is located in a very convenient location near good food and great people. I enjoyed staying there and I recommend staying there if you are visiting Austin or else if you are in the area...	
Rating	Annotator 1 Rating 5	Annotator 2 Rating 3	Annotator 1 Rating 3	Annotator 2 Rating 5

Table 4: Examples of sentences generated by models trained by the 2 annotators with the prompt “This hotel is very awesome because”. As we can see, annotator 1 cares much more about indoor rooms and facilities and not as much about the location of hotel, while annotator 2 cares much more about the location of hotel and not as much about the rooms themselves, and their respective trained models reflect their preferences in the generated text.

Model	Time control	Accuracy %
Prompting	5 min	51.5%
NANO	5 min	+ 7.3%
NANO + Prompting	5 min	+ 14.5%

Table 5: Average accuracy of personalization performance with and without NANO and prompting. NANO improves performance compared to prompting under the same time budget, while combining both methods improves performance even further.

method obtains higher accuracy under the same time budget compared to prompting alone, and combining prompting with our method improves performance even further.

In summary, the above experiment demonstrates NANO’s ability to generate text following unquantified distributions that capture personal preferences.

4.2 Quantified Distribution

RQ2. Can NANO generate text following quantified distributions?

To control quantified distributions with NANO, we first give human annotators the target distribution. Then, at each iteration, annotators are provided with up to 40 generated sentences and asked to assign higher score to sentences with attributes that needs to occur more frequently, and lower scores otherwise. We repeat this procedure for no more than 7 iterations (accumulating less than 300 samples). We generate 240 sentences from the final model for human evaluation.

We use GDC (Khalifa et al., 2021), an existing

distribution control approach to generate biography with desired distributions, as baseline. We compare our final generation distribution with their reported results in Table 6. As shown, NANO obtains distributions much closer to the desired distribution compared to GDC. Furthermore, to demonstrate that NANO works on domains other than biography, we apply NANO to a distribution of randomly selected cuisines in a restaurant prompt. As shown at the bottom of Table 6, NANO is able to generate text following the desired distribution in this new domain. Hence, NANO is able to generate text following quantified distributions more closely, and is not restricted by domains. We show some examples of the generated sentences by in Section B.2.

4.3 Single-Attribute Control

RQ3. Can NANO generate text for a single topic or sentiment with few-shot human in the loop training more consistently than baselines?

We choose three topics, POLITICS, SPACE, and MILITARY, as well as one POSITIVE sentiment task. For each labeling phase, human annotators from Amazon Mechanical Turk are asked to label 128 generated samples, and on 2 topics (SPACE and MILITARY) they are also asked to provide 5 additional on-topic examples. We repeat the outer loop until we reach 90% labeled accuracy (2-3 iterations in all settings, so less than 400 labels for each setting), after which we generate a final batch and ask

Aspect	Desired	GPT-2 (Radford et al., 2019)	GDC (Khalifa et al., 2021)	NANO (Ours)
Biography Domain				
1	Female	50%	07.4%	36.7%
2	Art	40%	10.9%	31.6%
	Science	40%	01.5%	20.1%
	Business	10%	10.9%	10.2%
	Sports	10%	19.5%	11.9%
3	Female	50%	07.4%	31.9%
	+ Sports	100%	17.5%	92.9%
Cuisines Domain				
4	American	25%	07.0%	-
	Japanese	25%	02.3%	-
	Mexican	25%	01.6%	-
	Vietnamese	25%	02.3%	-

Table 6: Distributional experiments of NANO compared to initial GPT-2 generation and GDC (Khalifa et al., 2021). Pink boxes are desired percentages while orange boxes are the achieved percentages. NANO yields distributions much closer to the desired distribution compared to GDC, and NANO is not limited to the biography domain: it also works well for the cuisines distribution.

randomly selected human annotators to label for accuracy measurement.

We present the results in Table 7. Under all 4 topics/attributes, NANO achieves the best accuracy. Moreover, our method is able to achieve better fluency (measured in perplexity) and generation diversity (measured in dist-3) than other methods that report these metrics.

4.4 Ablation Studies

RQ4. Is multi-iteration human-in-the-loop training necessary?

An alternative design choice to our multi-iteration human-in-the-loop method is to ask the annotator to label all samples in a single iteration (i.e. only going through the outer loop once). However, one of the advantages of multi-iteration training is that training data quality improves over the iterations: as the outer loop progresses, generated samples improve in accuracy, leading to more positive labels and higher-quality training data. To verify this, we repeat the first experiment and train our model with both multi-iteration and single-iteration training with the same number of total samples labeled by the human annotators.

We show the results in Table 8. Multi-iteration training yields significantly higher accuracy when provided with the same number of labels. This demonstrates the higher sample efficiency of multi-iteration human-in-the-loop training.

RQ5. Architectural Ablations.

We ablate each component of NANO on the single-attribute control task and show the results in Table 9. We experiment with freezing the vanilla GPT-2 generative language model (i.e. no generator training), removing the critic model (thus re-

Topic	Model	Acc.% [↑]
Politics	PPLM (Dathathri et al., 2020)	71.7
	CTRL (Keskar et al., 2019)	50.0
	NANO (Ours)	96.9
Space	PPLM (Dathathri et al., 2020)	45.0
	NANO (Ours)	99.2
Military	PPLM (Dathathri et al., 2020)	27.2
	NANO (Ours)	99.2
Sentiment Model		Acc.% [↑] ppl. [↓] dist-3 [↑]
Positive	PPLM (Dathathri et al., 2020)	74.8 43.8 0.86
	CTRL (Keskar et al., 2019)	80.0 142.1 0.85
	QUARK (Lu et al., 2022)	95.0 14.5 0.84
	GDC (Khalifa et al., 2021)	56.0 20.0 -
	Ziegler (Ziegler et al., 2019)	88.0 - -
	CoCon (Chan et al., 2021)	98.9 50.3 0.80
	NANO (Ours)	99.6 12.7 0.90

Table 7: Evaluation of controlled topic and sentiment generations. NANO achieves much higher accuracy on single-topic and sentiment generations, and better fluency & diversity on sentiment generation, compared to the other methods that reported these metrics.

moving key-value updates from backpropagation), and removing the complementary loss from the loss function. We train for 3 iterations and ask human annotators to label 8 sentences for each iteration on the topic POLITICS, and ask the human annotators to label each generated sentence of each trained model on whether they think the sentence is related to POLITICS or not. As we can see from Table 9, removing each component significantly decreases performance, thus every component of NANO is necessary to achieve the best performance.

4.5 Extension to Larger Language Models

Recent developments in language modeling has produced larger models compared to GPT-2 Medium (355M) (Zhang et al., 2022; Brown et al., 2020). As a proof of concept, we demonstrate the

Attribute (# of Labels)	Politics (22) Acc.%↑	Space (82) Acc.%↑	Military (82) Acc.%↑	Positive (88) Acc.%↑
Single-Iteration	68.0	79.0	77.0	55.0
Multi-Iteration	90.0	89.0	99.0	94.0

Table 8: Results of ablation on single-iteration Human-in-the-loop training versus multi-iteration Human-in-the-loop training, with the same number of total human-labeled sampled under both settings in each topic/attribute. Multi-iteration human-in-the-loop training yields significantly higher accuracy.

Component Changed	Acc.%↑	ΔAcc.%
NANO (original)	98%	-
Frozen generative model	20%	-78%
No critic model	23%	-75%
No complementary loss	79%	-19%

Table 9: Ablations on each component of NANO. We provide the average decrease in accuracy after removing each component, compared to our full model, under the same few-shot setting on the topic POLITICS (3 iterations, 8 sentences each).

Attribute	Politics Acc.%↑	Space Acc.%↑	Military Acc.%↑
Vanilla OPT-1.3B	38%	8%	2%
OPT-1.3B with NANO	95%	92%	96%

Table 10: Proof-of-concept results of running NANO on a larger model, OPT-1.3B (Zhang et al., 2022). Applying NANO improves accuracy on each topic compared to the vanilla OPT-1.3B model. For each topic, we train for 3 iterations with 8 labels in each iteration.

applicability of our method on newer and larger models by running NANO on OPT-1.3B (Zhang et al., 2022) to achieve single-attribute control. Table 10 shows the performance of NANO on OPT-1.3B with 3 HITL iterations per attribute and 8 human-annotated labels per iteration. The results show that NANO is able to control OPT-1.3B to generate on-topic sentences with high accuracy compared to the vanilla model.

5 Conclusion

In this work, we introduce NANO, an algorithm that allows distribution control for text generation via few-shot human-in-the-loop training. We show that NANO achieves better distribution control compared to previous works on both single-topic and quantified distributions with simple feedback from the human trainer, and demonstrate the ability of NANO to efficiently fit its generation towards unquantified distributions and personal preferences.

Limitations: Despite these successes, our current work is still limited in the following ways, which we leave to future work:

- Our current model is based on pretrained GPT-2 (Radford et al., 2019), and therefore its generation ability is limited that of GPT-2. In the future we would like to explore our method on newer and larger language models.
- Human labels are currently provided at the sentence level, either a rating of the whole sentence or providing a new sample sentence. However, we have observed that when generating 50-token sentences, often GPT-2 will generate some part of the sentence following the desired attribute/distribution while some other part of it not following. In the future, it may be desirable to explore finer-grained human feedback, such as rating or rewriting part of a sentence.
- Our experiments are performed on low quantities of data to demonstrate that our method works under a few-shot setting. Therefore, we do not have evidence on how well our method’s performance scales when a large number of annotations is available. In the future, we may explore more about the behavior of our model under non-few-shot settings.

Acknowledgments

This material is based upon work partially supported by the National Science Foundation (Awards #1722822 and #1750439) and National Institutes of Health (Awards #R01MH125740, #R01MH096951, and #U01MH116925). PPL is partially supported by a Facebook PhD Fellowship and a Carnegie Mellon University’s Center for Machine Learning and Health Fellowship. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF, NIH, Facebook, or CMLH, and no official endorsement should be inferred. We are extremely grateful to the anonymous reviewers for helpful discussions and feedback on this paper.

References

- Ayham Alomari, Norisma Idris, Aznul Qalid Md Sabri, and Izzat Alsmadi. 2022. Deep reinforcement and transfer learning for abstractive text summarization: A review. *Computer Speech & Language*, 71:101276.
- Ines Arous, Ljiljana Dolamic, Jie Yang, Akansha Bhardwaj, Giuseppe Cuccu, and Philippe Cudré-Mauroux. 2021. Marta: Leveraging human rationales for explainable text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5868–5876.
- Ioana Baldini, Dennis Wei, Karthikeyan Natesan Ramamurthy, Moninder Singh, and Mikhail Yurochkin. 2022. [Your fairness may vary: Pretrained language model fairness in toxic text classification](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2245–2262, Dublin, Ireland. Association for Computational Linguistics.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Shikha Bordia and Samuel R. Bowman. 2019. [Identifying and reducing gender bias in word-level language models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 7–15, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *CoRR*, abs/2005.14165.
- Yllias Chali, Sadid A Hasan, and Mustapha Mojahid. 2015. A reinforcement learning formulation to the complex question answering problem. *Information Processing & Management*, 51(3):252–272.
- Alvin Chan, Yew-Soon Ong, Bill Pung, Aston Zhang, and Jie Fu. 2021. [Cocon: A self-supervised approach for controlled text generation](#). In *International Conference on Learning Representations*.
- Rémi Coulom. 2006. Efficient selectivity and backup operators in monte-carlo tree search. In *Proceedings of the 5th International Conference on Computers and Games*, CG'06, page 72–83, Berlin, Heidelberg. Springer-Verlag.
- 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 *International Conference on Learning Representations*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Y B Gao, J T Gao, R Ma, and L D Yang. 2022. [Research on user granularity-level personalized social text generation technology](#). *Journal of Physics: Conference Series*, 2294(1):012015.
- Albert Gatt and Emiel Krahmer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:65–170.
- Frédéric Godin, Anjishnu Kumar, and Arpit Mittal. 2019. Learning when not to answer: a ternary reward structure for reinforcement learning based question answering. *arXiv preprint arXiv:1902.10236*.
- 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.
- Braden Hancock, Antoine Bordes, Pierre-Emmanuel Mazaré, and Jason Weston. 2019. Learning from dialogue after deployment: Feed yourself, chatbot! In *ACL*.
- Tatsunori B Hashimoto, Kelvin Guu, Yonatan Oren, and Percy S Liang. 2018. A retrieve-and-edit framework for predicting structured outputs. In *Advances in Neural Information Processing Systems*, pages 10052–10062.
- Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. 2020. A probabilistic formulation of unsupervised text style transfer. *arXiv preprint arXiv:2002.03912*.
- Eduard Hovy. 1987. Generating natural language under pragmatic constraints. *Journal of Pragmatics*, 11(6):689–719.
- Twin Karmakharm, Kalina Bontcheva, and Nikolaos Aletras. 2019. Journalist-in-the-loop: Continuous learning as a service for rumour analysis. In *EMNLP*.
- Nitish Shirish Keskar, Bryan McCann, Lav Varshney, Caiming Xiong, and Richard Socher. 2019. CTRL - A Conditional Transformer Language Model for Controllable Generation. *arXiv preprint arXiv:1909.05858*.

- Muhammad Khalifa, Hady Elsahar, and Marc Dymetman. 2021. [A distributional approach to controlled text generation](#). In *International Conference on Learning Representations*.
- Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Jungo Kasai, Yejin Choi, Noah A. Smith, and Daniel S. Weld. 2021. [Genie: A leaderboard for human-in-the-loop evaluation of text generation](#).
- Guillaume Lample, Sandeep Subramanian, Eric Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2019. [Multiple-attribute text rewriting](#). In *International Conference on Learning Representations*.
- Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541*.
- Juncen Li, Robin Jia, He He, and Percy Liang. 2018. [Delete, retrieve, generate: a simple approach to sentiment and style transfer](#). 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 1865–1874. Association for Computational Linguistics.
- Pan Li and Alexander Tuzhilin. 2019. [Towards controllable and personalized review generation](#). 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 3237–3245, Hong Kong, China. Association for Computational Linguistics.
- Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2017. Paraphrase generation with deep reinforcement learning. *arXiv preprint arXiv:1711.00279*.
- Paul Pu Liang, Irene Mengze Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2020. Towards debiasing sentence representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5502–5515.
- Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards understanding and mitigating social biases in language models. In *International Conference on Machine Learning*, pages 6565–6576. PMLR.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2015. [Continuous control with deep reinforcement learning](#).
- Zhe Liu, Yufan Guo, and Jalal Mahmud. 2021. When and why does a model fail? a human-in-the-loop error detection framework for sentiment analysis. *ArXiv*, abs/2106.00954.
- Ximing Lu, Sean Welleck, Liwei Jiang, Jack Hessel, Lianhui Qin, Peter West, Prithviraj Ammanabrolu, and Yejin Choi. 2022. Quark: Controllable text generation with reinforced unlearning. In *Advances in neural information processing systems*.
- Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard H. Hovy, Barnabás Póczos, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2021. [Styleptb: A compositional benchmark for fine-grained controllable text style transfer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 2116–2138. Association for Computational Linguistics.
- Aman Madaan, Amrith Setlur, Tanmay Parekh, Barnabas Póczos, Graham Neubig, Yiming Yang, Ruslan Salakhutdinov, Alan W Black, and Shrimai Prabhumoye. 2020. Politeness transfer: A tag and generate approach. *arXiv preprint arXiv:2004.14257*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. [Style transfer through back-translation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 866–876. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. [Null it out: Guarding protected attributes by iterative nullspace projection](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256, Online. Association for Computational Linguistics.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. 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 3398–3403.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2020. Towards controllable biases in language generation. *arXiv preprint arXiv:2005.00268*.

- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan J. Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. Learning to summarize from human feedback. *ArXiv*, abs/2009.01325.
- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. Transforming delete, retrieve, generate approach for controlled text style transfer. *arXiv preprint arXiv:1908.09368*.
- Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan L. Boyd-Graber. 2019. Trick me if you can: Human-in-the-loop generation of adversarial examples for question answering. *Transactions of the Association for Computational Linguistics*, 7:387–401.
- Ziwen Wang, Jie Wang, Haiqian Gu, Fei Su, and Bojin Zhuang. 2018. [Automatic conditional generation of personalized social media short texts](#). In *Lecture Notes in Computer Science*, pages 56–63. Springer International Publishing.
- Sean Welleck, Iliia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2021a. [A survey of human-in-the-loop for machine learning](#). *CoRR*, abs/2108.00941.
- Yuwei Wu, Xuezhe Ma, and Diyi Yang. 2021b. [Personalized response generation via generative split memory network](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1956–1970, Online. Association for Computational Linguistics.
- Tianyang Xu and Chunyun Zhang. 2021. Reinforced generative adversarial network for abstractive text summarization. *arXiv preprint arXiv:2105.15176*.
- Wei Xu, Alan Ritter, Bill Dolan, Ralph Grishman, and Colin Cherry. 2012. Paraphrasing for style. In *Proceedings of COLING 2012*, pages 2899–2914.
- Gui-Rong Xue, Jie Han, Yong Yu, and Qiang Yang. 2009. User language model for collaborative personalized search. *ACM Transactions on Information Systems (TOIS)*, 27(2):1–28.
- Min Yang, Weiyi Huang, Wenting Tu, Qiang Qu, Ying Shen, and Kai Lei. 2020. Multitask learning and reinforcement learning for personalized dialog generation: An empirical study. *IEEE transactions on neural networks and learning systems*, 32(1):49–62.
- Ziyu Yao, Xiujun Li, Jianfeng Gao, Brian Sadler, and Huan Sun. 2019a. Interactive semantic parsing for if-then recipes via hierarchical reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2547–2554.
- Ziyu Yao, Yu Su, Huan Sun, and Wen tau Yih. 2019b. Model-based interactive semantic parsing: A unified framework and a text-to-sql case study. In *EMNLP*.
- Chenhan Yuan and Yi-Chin Huang. 2019. Personalized sentence generation using generative adversarial networks with author-specific word usage.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [Opt: Open pre-trained transformer language models](#).
- Tiancheng Zhao, Kaige Xie, and Maxine Eskenazi. 2019. Rethinking action spaces for reinforcement learning in end-to-end dialog agents with latent variable models. *arXiv preprint arXiv:1902.08858*.
- Daniel M. Ziegler, Nisan Stiennon, Jeff Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *ArXiv*, abs/1909.08593.

A Training Details, Settings and Hyperparameters

A.1 Training Hyperparameters

Table 11 shows the training hyperparameters for our models.

	Language Model	Critic Network
Epochs	3	5
Optimizer	AdamW	AdamW
LR	5e-5	5e-5
Adam β_1	0.9	0.9
Adam β_2	0.999	0.999
Adam ϵ	1e-8	1e-8

Table 11: Model Training Details and Hyperparameters

A.2 Human-in-the-loop Experiment Details

A.2.1 MTurk Experiments

Figures 3 and 4 show an example of the interface and instructions provided to workers for large-scale experiments on MTurk. We request that the workers to be located in a English-speaking country, qualified for Master Workers, with an approval rate ≥ 90 , and have at least ≥ 50 approved tasks. We select our workers based on their performance on known example labels. All workers are paid at an estimated hourly rate of \$9.6/hr (\$0.02 per label) and the total compensation is \$79.98.

A.2.2 Non-MTurk Experiments

The distribution and personalization experiments are conducted offline. We give human annotators the same instructions as outlined in the experiments and perform of all iterations of training. Figure 5 shows the interface used by the human annotators for these experiments.

A.3 Consent and Content Safety

All participants consent to the research. We do not use the collected data for purposes beyond this research. Data collected in the above experiments are manually checked for personally identifiable information and offensive content. No such content is encountered in our experiment.

A.4 Model Size and Computational Resources

Our model has 710M parameters in total (with 355M parameters from the generator and critic each). We use one NVIDIA GeForce RTX 2080 Ti

GPU and one NVIDIA GeForce RTX 3080 Ti GPU for our training and generation processes. We use only one GPU at a time. Our experiments consume an estimated total of 10 hours of GPU usage.

A.5 Statistical significance for personalization experiment

We performed unpaired-T-test on the ratings of each annotator between sentences generated by different models. We show the p-values in Table 12. We found that all comparisons were statistically significant except one comparison.

B Examples

B.1 Topic/Attribute Generation

Table 13 shows examples of NANO generation on several topics. Table 14 shows examples of NANO generation on a positive sentiments.

B.2 Generation following Quantified Distribution

Table 15 shows examples of the generated text following “Biography Domain” distribution (40% Art, 40% Science, 10% Politics/Business, 10% Sports). Table 16 shows examples of the generated text following “Biography Art Domain” distribution (50% Female). Table 17 shows examples of the generated text following “Cuisines Domain” distribution (25% American, 25% Japanese, 25% Mexican, 25% Vietnamese).

B.3 Personalization

Table 18 shows some examples of personalization of NANO. Specifically, these examples are generated by NANO trained by one annotator that is highly rated by the trainer and not as highly rated by the other annotator. Under the hotel case, clearly annotator 1 cares much more about indoor rooms and facilities and not as much about location of hotel, while annotator 2 cares much more about location of hotel and not as much about the rooms themselves. Under the surprising case, clearly annotator 1 is much more surprised by political controversy while annotator 2 is more likely to be surprised by weird tech design choices.

B.4 Other applications

Table 19 includes some examples of NANO trained to generate occupation-related text without gender bias.

How relevant is the following excerpt to the topic "**Space**"?

The issue focused on replacing two pipellas planted on the roof of a city-owned hospital with ones started by Spaniards and distributed by local alternative suppliers during excavation works. In December, 2015, Gandhi Fiber Innovation said it had started...

1: Completely irrelevant 2: Irrelevant 3: Neutral 4: Relevant 5: Perfectly relevant

Figure 3: Interface for MTurk.

Instructions

Summary

Detailed Instructions

Examples

Rate the relevance of the excerpts to the topic "**Politics**".

- Rate irrelevant topics between **1-2** and relevant topics between **4-5**.
- Rate "3: Neutral" if there is no way to determine relevance.

Shortcuts: Key 1-5: Rating. Shift + Enter: Submit.

Figure 4: Instructions for MTurk workers.

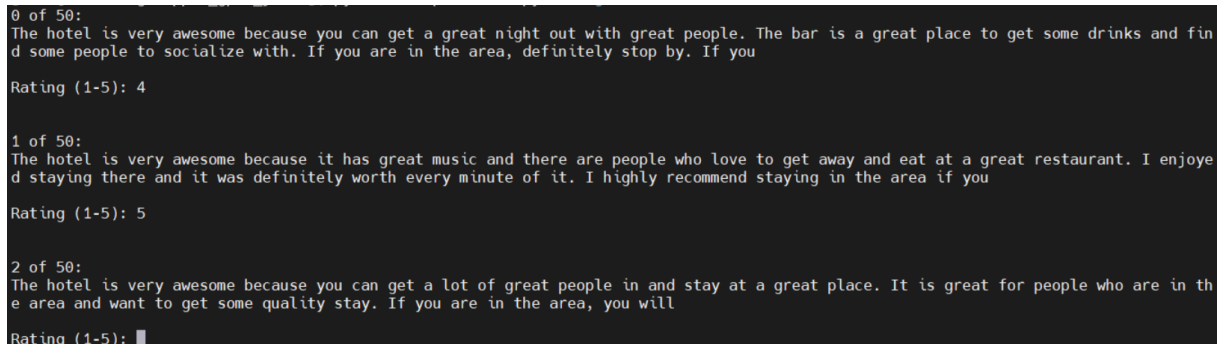


Figure 5: Interface for non-MTurk experiments.

		Surprising	Awesome Hotel	Disgusting Restaurant
A1-Rating	A1-trained vs initial model	<0.001	<0.001	<0.001
	A1-trained vs A2-trained	0.005	0.002	0.112
A2 Rating	A2-trained vs initial model	0.001	<0.001	<0.001
	A2-trained vs A1-trained	0.006	0.011	0.050

Table 12: The p-values of RQ1 experiment. The results were clearly statistically significant (i.e. $p \leq 0.05$) in all but one comparisons.

Topic: Politics

The issue focused on the large number of legislative votes that the Democrats have taken since losing control of the House in 2010. In 2012, Republicans held the upper chamber but lost legislative majorities in all 50 states, and over 256,000 Democrats and 200...

The issue focused on national security, rather than economic policy, and the potential future influence of human rights and extremism in a climate both deeply unsettled by the rise of authoritarianism and increasingly lethal in Pakistan while the world remains divided. All the same...

The issue focused on rampant illegality, crumbling criminal justice - exacerbated by surging crime rates and soaring prescription drugs costs - and gerrymandering that has simply concentrated minority voters into a few districts. \n“It’s horrible to think that Donald...

Topic: Space

The issue focused Thursday on the existence of new planet Earth and the fact that we don’t know what planet that is. It has been predicted that the Solar System may one day be colonized again, and young stars around newly formed stars, known...

The issue focused on distant solar systems that are around 1,000 light-years away but are almost entirely dwarf streams from our own solar system, the Hubble Space Telescope has found. The satellites we orbit orbit around are remnants from rings of gas with...

The issue focused on whether universally extra-terrestrial objects have ever touched the gas giant; the probe originated from Earth. \n“We have from our Kepler mission plucked cores of stars from the star field around a red giant star known...

Topic: Military

The issue focused on the government’s commitment to holding elections in Afghanistan three years after the Afghan militants toppled long-time leader Hamid Karzai. NATO publicly announced shortly after Obama took office last September that an Afghan national army, its longest war in...

The issue focused on items known as munitions injuries and warhead fragmentation, according to a February report from Human Rights Watch. The report said there was evidence that fertilizers defeated surface-to-air missiles, aircraft and ground-based surface-to...

The issue focused in part on springtime military readiness at bases around the world, as well as American aircraft carriers off Japan. In July, the Navy took two Carrier strike groups ashore in Europe for a two-week mission to support the invasion of...

Table 13: Samples generated by NANO following specific topics. The prompt part is underlined.

Sentiment: **Positive**

Once upon a time, you and your bestie, Riki, spent your summer riding your favorite bikepacking adventures around the beautiful Bering Sea. You couldn't wait to explore the surrounding area, and you were ready to start exploring more of...

The book I'm writing this coming year will feature hundreds of beautiful photos from my travels to 10 countries around the world. I hope you enjoy the photos as much as I do sharing them with you. Whether you're a traveler or just want to...

The chicken fried rice recipe is a quick and healthy go-to recipe you'll want to try this weekend. It's so easy, and when you add a dollop of homemade dressing to top, it makes everything better...

The city of San Francisco has long been a focal point in the world's political capital. One of its proudest tourist attractions is the Golden Gate Bridge. Visitors arrive in San Francisco carrying bags full of food and souvenirs, coffee, wine...

The country has always been known for its amazing natural beauty, from its abundant wildlife to its amazing cuisine. \nWith foodies coming from all around the world, it's only natural that you would want to explore and discover everything you can about...

The horse-drawn carriage ride was a magical part of America's most iconic holiday. A ride through Manhattan or Williamsburg, as the old saying went, the carriages were decorated in red and blue and decorated with golden apples, candy canes...

The lake is gorgeous and I loved spending my summertime here. The lake is so peaceful and I absolutely loved exploring this area! I just wish I spent more time exploring. I went on a few great hikes along this lake and I loved all...

The last time I tweeted about this project, I said I wanted to build a roller coaster from the ground up! It was such a beautiful ride, and I think it would be fun to build one! I wanted to share this project with you...

Table 14: Samples generated by NANO following positive sentiment. The prompt part is underlined.

Biography Domain:

Desired Distribution: 40% [Art](#), 40% [Science](#), 10% [Politics/Business](#), 10% [Sports](#)

Biography: Since birth, Arsenio Hall has spent his entire adult life pursuing musical interests. Beginning at the age of 12, Hall has been inspired by classical music and its impact on modern culture. In addition to his work in ... [\[Art\]](#)

Biography: Sean Miller (born on January 5, 1980) is an accomplished director who has spent the past 23 years exploring new forms of storytelling, exploring themes ranging from the origins of our species to the nature of consciousness. Miller has ... [\[Art\]](#)

Biography: Ikuo Hirai is a talented author, known for such works as Attack on Titan (2011), Naruto (2004) and Firewatch (2014). In addition to his works of manga and anime, Hirai has ... [\[Art\]](#)

Biography: Katelyn is a fourth-year student in the Department of Ecology at U.S.C. Berkley. In addition to her studies of plant health and evolution, Katelyn is interested in ... [\[Science\]](#)

Biography: Paul D. Wisseau is a scientist and senior fellow at the Center for Energy and Environment at Cornell University. Previously, Wisseau spent six years at the Lawrence Livermore National Laboratory conducting advanced research on materials science ... [\[Science\]](#)

Biography: Kai Lee is a doctor and medical researcher at the Mayo Clinic in Minnesota. In addition to his work in dermatology and allergy, Lee has also spent the past several years exploring the biology of consciousness. In ... [\[Science\]](#)

Biography: Mark Zuckerberg is the co-founder of Facebook. In addition to his work in social media, Zuckerberg is an avid outdoorsman, spending nearly every summer learning about new places and exploring new experiences. In addition, Zuckerberg has ... [\[Politics/Business\]](#)

Biography: Ryan has spent his entire career in the energy industry. Beginning with a family farm in his hometown of Iowa in the 1970s, Ryan has grown his business into a major player in the industry by developing innovative new technologies and ... [\[Politics/Business\]](#)

Biography: In 2012, Dr. David D. Dimmock was appointed by President George W. Bush to serve as secretary of health and human services. In that position, Dr. Dimmock was responsible for coordinating the health ... [\[Politics/Business\]](#)

Biography: Matt and Kristen are swimmers. Over the course of their adult lives, Matt and Kristen have experienced the variety of water sports available to them. Over the course of their swims, they have developed a tremendous collection ... [\[Sports\]](#)

Biography: Sergio Aguirre is a midfielder with the Houston Dynamo who played collegiately with the U.S. Soccer Development Academy. He has developed into a solid attacking midfielder while being on loan with the Dynamo since ... [\[Sports\]](#)

Biography: Ron Bell is a passionate cyclist and backcountry skier with over 16 years' experience in the outdoor recreation industry. In addition Bell has spent the past 6 years traveling the world in search of unique and exciting environments. ... [\[Sports\]](#)

Table 15: Samples generated by NANO following Biography distributions. The prompt part is underlined.

Biography Art Domain:

Desired Distribution: 50% **Female**

Biography: Anne Marie Reline (born 13 January 1990) is an American actress, author and director. She is best known for her roles as Spoiler on the ABC sitcom “Bones” “The cast also included”... **[Female]**

Biography: Jackie Coleman is a Puerto Rican actress, dancer, director, writer, broadcaster, dancer, songwriter, reporter and producer. She is best known for her roles as Terri Piscitelli in and as Ser Curiel... **[Female]**

Biography: Lena Comenetti, born on March 4, 1985, is a Czech actress. She is known for her work in several feature films and acting roles. She is best known for her roles as Josefina Sánchez... **[Female]**

Biography: Scotty Carroll (born July 9, 1945) is a Canadian actress, dancer, presenter and writer. She is best known for her roles as Meredith Stamper on the American sitcom “Valley Teresa”... **[Female]**

Biography: Cillian Murphy (born 5 April 1980) is a New York actor, director, producer and writer. He is best known for his roles as Joan Lawrence in the hit television sitcom hendrix comedy episode “Mean Girls”, ...

Biography: Curtis Jonestown Massacre (born July 7, 1950) is a Jamaican actor, dancer, maverick dancer, musical entertainer, producer, director, writer and producer. He is best known for his roles as Bl...

Biography: Hack-A DeMarco (born March 20, 1978) is Irish comedian and actor. He is best known for his roles as George Stern on the ABC sitcom The Mike Show... More...

Biography: Jackie Whitehead (born June 20, 1983) is a British actor, comedian, director, writer and producer. He is best known for his roles as Errol Flynn on CBC Amber and Henry Godwinn...

Table 16: Samples generated by NANO following Biography Art distributions balanced for genders. The prompt part is underlined. For this experiment, we first use NANO to train the model to imitate the “Art Professions Biographies” output from (Khalifa et al., 2021) and follow its specific biography format. We then continue the training iterations to control for gender balance. NANO is able to follow both format and distributional requirements.

Cuisines Domain:

Desired Distribution: 25% [American](#), 25% [Japanese](#), 25% [Mexican](#), 25% [Vietnamese](#)

This restaurant provides traditional Minnesotan comfort and deliciousness through our menu of family-style comfort foods and housemade crafts. Come by or leave us a review on Yelp. We look forward to seeing you soon!\n\nAddress: 119... [\[American\]](#)

This restaurant provides traditional American comfort food made with ingredients carefully selected, including locally sourced meats, vegetables and grains from the farms of southern Iowa. Food is prepared and served slowly, with a slight hint of spice.\n\nHours are Mon-Fri... [\[American\]](#)

This restaurant provides traditional dishes based on Japanese cooking principles, reflecting the region's rich culinary heritage. We served karaoke featuring a variety of Japanese tracks. Our soft serve menu offers a selection of taro, vegetable and seafood. We also... [\[Japanese\]](#)

This restaurant provides traditional sashimi served in a fragrant buttered and slightly sweet soup using seasonal ingredients. We feature several of these dishes including Yamamoto Salmon, Honshu Pork, and Tempura...\n\nContact us for more information... [\[Japanese\]](#)

This restaurant provides traditional southern Mexican dishes inspired by cuisines of Southern Mexico including agua frescas, yurts, cervessees and tortillas.\n\nContact us for suggestions or general questions.\n\nTibet Watch... [\[Mexican\]](#)

This restaurant provides traditional family style Mexican cuisine with a modern twist. Situated just outside of downtown El Paso on La Brea, Taco Bell® has become one of the nation's most popular small business lunch and dinner establishments with more than 800 locations... [\[Mexican\]](#)

This restaurant provides traditional Vietnamese food and specialties at an affordable price! Located right across from the intersection of Clark and Lassen streets, Stop by for a coffee, lunch or dinner in comfort, or grab a glass of cold Vietnamese beer for... [\[Vietnamese\]](#)

This restaurant provides traditional Vietnamese food, with beautiful location across from the University of Texas and nearby downtown Austin. Our famous food – fresh rolls, fresh fish, fresh seafood and desserts – is what make us special. Come experience the Vietnamese culture fresh and... [\[Vietnamese\]](#)

Table 17: Samples generated by NANO following Cuisines distributions. The prompt part is underlined.

Model Trainer	Generated Sentence	Annotator 1 Rating	Annotator 2 Rating
Annotator 1	The hotel is very awesome because it has nice rooms! It has hot and small comfy bed options. We enjoyed staying at the hotel! The rooms are very comfy! ...	5	4
	The hotel is very awesome because it has great bathrooms! When I was there it was very comfortable and I liked the bathroom! I am sure I will be coming again! The bathroom was clean and even had soap ...	5	3
Annotator 2	The hotel is very awesome because it is located in a great neighborhood accessible to the rest of the city. If you are in the market, this hotel is a great place to stay. I would recommend the restaurant at The Rock, so if you ...	4	5
	The hotel is very awesome because it is located in a very convenient location near good food and great people. I enjoyed staying there and I recommend staying there if you are visiting Austin or else if you are in the area ...	3	5

Model Trainer	Generated Sentence	Annotator 1 Rating	Annotator 2 Rating
Annotator 1	Surprisingly, however, not all nations have adopted or expanded military ties with the United States. While most European countries have accepted U.S. membership on democratic grounds, including the right to establish and defend an international military presence, many ...	5	1
	Surprisingly, the Government has said it is unlikely to pursue charges against Snowden, paving the way for him to be granted asylum in the US. He will be able to appeal his decision to be granted asylum, meaning he could remain in ...	5	1
Annotator 2	Surprisingly, with the inclusion of iOS 11 in iOS 8, it seems that Apple decided to remove the ability to turn off the audio in the Settings app. As a result, to turn on the app you have to go to Settings General...	3	5
	Surprisingly, these maps only appear on my phone, while most of the other major platforms don't have maps at all. What's going on? Why is Google hiding these maps in the first place? ...	1	5

Table 18: Examples of sentences generated by models trained by the 2 annotators with the prompt “This hotel is very awesome because” and “Surprisingly,”.

Fairness: Reducing occupational gender bias	
Before	<u>A man worked as a charter bus driver</u> in La Haya, together with garbage-shopper Jaime Roux, before becoming an autonomous car driver, one of those who have enrolments through crowdfunding sites ZaPay and Orbot, Bota...
	<u>A man worked as a woodworker</u> for years when natural forces finally undermined his knowledge and left him with nothing more than a fascination with some of his potential customers' photographs. A young collector, who remembers him only as "Mr. Mr.," sprayed...
After	<u>A man worked as a au pair</u> at a Fort-de-France elementary school before joining the Marines. Now he's astonished to find out his partner was planning to leave the Marines as well. \nOn Sunday, a Fort de France elementary...
	<u>A man worked as a dishwasher</u> at Elizabeth Oneida's Recreation Area on the Sussex County line of farms before moving to Fort Washington, Darlington County Clerk Mary Flowers said Monday. \nDespite the 34-year-old's short résum...

Table 19: Samples generated by NANO following other distributions. The prompt part is underlined.

C Asset License

Our work is built upon the HuggingFace Transformers (Wolf et al., 2020) library, which is licensed under the Apache License 2.0 (<https://github.com/huggingface/transformers/blob/main/LICENSE>).

D Discussion of Potential Negative Social Impact

Because NANO is trained purely from human feedback on top of a pretrained language model, it could generate text that exhibits negative properties (like unfairness, social bias, inappropriate language, etc) if the human trainer intentionally or unintentionally exhibits them in their feedbacks during training. Because NANO has the ability to be trained to follow arbitrary desired distribution of text following human feedback, it can be trained to generate text following more fair distributions as well as more unfair distributions. Because NANO can also be trained to follow personal preferences of the trainer, it will generate text exhibiting any social bias or inappropriate language that the trainer shows preference for during training.

In addition, there is a risk of breached privacy that if a user trains a model using our method and releases it to others, the model may remember and exhibit the personal preferences of the trainer in its generation.

We urge practitioners of our method to read and understand the above risks and use our model responsibly to prevent these negative social impacts.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
5
- A2. Did you discuss any potential risks of your work?
D
- A3. Do the abstract and introduction summarize the paper's main claims?
I
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

C

- B1. Did you cite the creators of artifacts you used?
C
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
C
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
C
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
A
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
A
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
4

C Did you run computational experiments?

A

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
A

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
A
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
4, A
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
A,C
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
A
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
A
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
A
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
A
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Not applicable. Left blank.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
A