A STEP-WISE WEIGHTING APPROACH FOR CONTROLLABLE TEXT GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We study the problem of controllable text generation (CTG): steering a language model (LM) to generate text with a desired attribute. Many existing approaches either require extensive training/fine-tuning of the LM for *each* single attribute under control or are slow to generate text. To this end, we first propose a framework based on step-wise energy-based models (EBMs) that is efficient in sampling and flexible in a wide range of practical CTG scenarios. Indeed, a number of existing CTG methods are special instances of our framework with a specific EBM design. In different control scenarios, we then design the respective energy functions that strategically up- or down-weigh the probabilities of keywords associated with a certain control attribute at each generation step. In experiments, we show that our simple and efficient approach is surprisingly competitive against more computationally expensive strong baselines, and even achieving new state-of-the-art performances in several cases. Our framework also provides a tuning hyper-parameter that nicely trades off generation quality and control satisfaction, enabling practitioners to easily adjust it to meet their needs.

1 INTRODUCTION

Large language models (LMs) pre-trained on web-scale data have now become an integral part of many natural language processing (NLP) applications. However, most pre-trained LMs can only perform generic generation that is unconstrained: after users provide the LM a piece of text as the input prompt, they have little control over the generative process (Holtzman et al., 2020). In many real-world cases, users may wish to impose some constraints or controls that guide the LMs towards generating texts of their specified attributes on the fly. Interesting applications include poetry generation with specific rhythms and topics (Wang et al., 2016), story generation with a pre-specified story line (Peng et al., 2018), and product review generation with a particular sentiment (Zang & Wan, 2017). Moreover, the ability to impose certain controls over the text generation process is critical for the safe and responsible deployment of LMs in practice. For example, prior research has shown that LMs may resemble or even amplify human biases (Sheng et al., 2019; Bordia & Bowman, 2019) and that LMs can be easily triggered to generated toxic content (Bender et al., 2021; Gehman et al., 2020). Both biased and toxic generated text should be avoided in most practical use cases.

In this paper, we focus on the problem of *controllable text generation* (CTG) that underlies the above applications. Given one or more desired, user-specified control attribute(s), such as "positive sentiment" and "science topic", we would like to steer an LM to generate text that satisfies the control attribute(s). As an example, Table 1 illustrates the CTG problem with various control attributes. Recently, a number of papers have proposed different methods to tackle this problem but they suffer from many practical issues. For example, some of these methods either require extensive training/fine-tuning for each CTG scenario (Khalifa et al., 2021; Bakhtin et al., 2021; Bhattacharyya et al., 2021) or are computationally expensive during generation (Liu et al., 2021; Dathathri et al., 2020; Shirish Keskar et al., 2019). Some methods (Ross et al., 2021; Liang et al., 2021; Ghazvininejad et al., 2017) are only applicable to one or two selected CTG tasks. Our goal in this paper is to develop a CTG method that is more efficient in sampling and more applicable to various CTG scenarios than the previous methods, without sacrificing the overall performance.

We propose a framework based on a *step-wise energy-based model* (EBM) for CTG. Our framework leverages the fact that modern LMs generate texts in an auto-regressive manner: they generate one word (token) at a time based on the previously generated words (tokens). We thus formulate the

Table 1: Illustration of different texts generated from the original GPT-2 and our CTG method, respectively, which demonstrates that our method can successfully steer GPT-2 toward distinct desired attributes.

GPT-2, unconstrained: As we approach the year 2023, one of the major problems that is facing society is the lack of information about the history of the planet, the importance of the planet to the world, and	
GPT-2 + word "restaurant"	**GPT-2 + topic "space"**
As we approach the year 2023, one of the major restaurant brands in America is looking at the potential of the future of food and how food should be made. I've been involved in	As we approach the year 2023, starship design is still a hot topic in the galaxy, as well as the future of space travel. In the latest issue of Wired magazine, Chris Roberts talks a
GPT-2 + positive sentiment	**GPT-2 + negative sentiment**
You can watch us perform live in the real world, right here! You'll see us perform live at DreamWorks' home in Los Angeles, complete with live artworks by the likes of DreamWorks	As I leave for the end of the week, I have a few questions for anyone who is considering leaving this blog. My name is John. I'm unemployed and run a small blog called

CTG problem as steering the word distribution at each generation step via a specially designed energy function. Our framework is general: we show that many existing CTG approaches are in fact special instances of our framework. Under our framework, we also develop a generic form of energy functions that can be applied to a wide range of CTG tasks. Intuitively, the energy function imposes a control of the generated text by adjusting the weights of keywords associated with the desired attribute at each generation step.

Via a range of experiments, we show the wide applicability of our CTG method by instantiating it for five commonly studied CTG tasks including i) word control, ii) topic control, iii) sentiment control, iv) text detoxification, and v) compositions of multiple controls. Our method is highly effective: on tasks iii) and iv), our method achieves competitive results compared to strong baselines. On tasks i), ii), and v), our method establishes new state-of-the-art CTG performances. Our method is also highly computationally efficient: it requires no training/fine-tuning, uses up to 18× fewer parameters, and is up to 100× faster than the baselines. We summarize our contributions as follows.

[C1] We propose a new CTG framework based on the step-wise EBM that weights the sampling distribution at each generation step through the energy function. Our framework unifies a number of prior approaches and enables developing new ones via specific energy function designs (Section 2.2).

[C2] We develop a generic form of an energy function for our framework. Our energy function is simple but powerful and flexible, which can be instantiated in different ways in a wide range of CTG scenarios with infinitesimal computational overhead (Section 2.3).

[C3] We validate the effectiveness of our method on five CTG tasks. Extensive experiments show that our method achieves highly competitive performance, often outperforming strong baselines, and is the most efficient method among all baselines (Section 3).

2 METHOD

In this section, we describe our approach in detail. We first review language models and motivate our approach with the sequence-level formulation of the CTG problem. We then introduce our step-wise EBM framework and propose an efficient and flexible form of energy function for a variety of CTG tasks. We close this section by reviewing related literature and, in particular, showing how our framework unifies a number of prior CTG methods.

2.1 PRELIMINARIES

Language models (LMs). Let $x \in \mathcal{V}$ be a word token, i.e., the smallest element we consider in modeling language, from a pre-specified token set $\mathcal{V}$ that we denote as the "vocabulary". Let $\boldsymbol{x} = \{x_1, \ldots x_T\} \in \mathcal{X}$ be a sequence of T tokens representing a piece of text, where $\mathcal{X}$ denotes the set of all such texts. An LM p_θ then defines a probability mass function on $\mathcal{X}$. In this paper, we work with *auto-regressive* LMs, namely,

$$p_\theta(\boldsymbol{x}) = \prod_{t=1}^{T} p_\theta(x_t | x_{1:t-1}), \quad \text{where } x_{1:t-1} = \{x_1, \ldots, x_{t-1}\}. \tag{1}$$

In particular, we denote $x_{1:0} = \emptyset$ such that $p_\theta(x_1|x_{1:0}) := p(x_1)$. Note that the dominant LMs, such as GPTs (Radford et al., 2019; Brown et al., 2020) and Recurrent Neural Networks (Mikolov et al., 2010; 2011), are formulated in this way.

Sequence-level EBM for CTG. The CTG problem can be formulated as learning a joint distribution $p_\theta(\boldsymbol{x}, c)$ of the text $\boldsymbol{x}$ and the control attribute c, which we define as

$$p_\theta(\boldsymbol{x}, c) = p_\theta(\boldsymbol{x}) p_\theta(c|\boldsymbol{x}) \propto p_\theta(\boldsymbol{x}) e^{-E_\theta(c|\boldsymbol{x})}\,, \tag{2}$$

where $p_\theta(\boldsymbol{x})$ is the base LM (such as GPT) and $p_\theta(c|\boldsymbol{x})$ is a conditional distribution of attribute given a sequence $\boldsymbol{x}$, modeled with the conditional energy function $E_\theta(c|\boldsymbol{x})$. Intuitively, low energies indicate more control satisfaction; thus, this model steers the base LM $p_\theta(\boldsymbol{x})$ towards generating text samples $\boldsymbol{x}$ with low energies.

The formulation in Eq. 2 is known as the *joint energy-based model* (EBM), which has been widely explored in the vision domain (Grathwohl et al., 2020; Che et al., 2020). More recently, Khalifa et al. (2021) adopted this joint EBM formalism for a sequence-level sampling of texts and proved that Eq. 2 is the optimal solution to CTG as a constraint optimization problem. We thus term their formulation as the *sequence-level EBM*. Notably, Eq. 2 also enables easy compositionality. If we have multiple attributes $\{c_i, \cdots, c_N\}$ to control and if we assume conditional independence of the their respective energy functions $E_\theta(c_i|\boldsymbol{x})$, then we can simply compute the resulting joint distribution by the summation of energies $\sum_{i=1}^N E_\theta(c_i|\boldsymbol{x})$ in the exponential term via

$$p_\theta(\boldsymbol{x}, c_1, \ldots, c_N) = p_\theta(\boldsymbol{x}) \prod_{i=1}^{N} p_\theta(c_i|\boldsymbol{x}) \propto p_\theta(\boldsymbol{x}) e^{-\sum_{i=1}^N E_\theta(c_i|\boldsymbol{x})}\,. \tag{3}$$

2.2 From Sequence-level EBM to Step-Wise EBM

Despite its appealing properties, such as theoretical optimality and compositionality, sequence-level EBM suffers from computational inefficiency that hinders its practical usage. Indeed, sampling from the joint distribution $p_\theta(\boldsymbol{x}, c)$ poses a significant challenge because of the huge sampling space with d^T possible sequences where d and T denote the vocabulary size and sequence length. It becomes more challenging to generate longer sequences with exponentially large sampling time.

We now introduce our step-wise weighting approach to approximate the expensive sampling step in the sequence-level EBM. Our approach leverages the nature of auto-regressive LM $p_\theta(\boldsymbol{x})$ that enables efficient sampling x_t at each time step t while satisfying the attribute constraint. To see this, we first decompose $p_\theta(\boldsymbol{x}, c)$ as

$$p_\theta(\boldsymbol{x}, c) = p(c) \prod_{t=1}^{T} p_\theta(x_t|c, x_{1:t-1})\,, \tag{4}$$

where, using Bayes' Rule, each term inside the product on the right-hand side becomes

$$p_\theta(x_t|c, x_{1:t-1}) = \frac{1}{Z} p_\theta(x_t, c|x_{1:t-1})\,, \tag{5}$$

where $Z := p_\theta(c|x_{1:t-1})$ is a constant due to the *independence assumption* which states that the previously generated text $x_{1:t-1}$ and the control attribute c are independent at each step t during generation, i.e., $p_\theta(c|x_{1:t-1}) = p(c)$. Leveraging this assumption and plugging Eq. 5 into Eq. 4, we have

$$\begin{aligned} p_\theta(\boldsymbol{x}, c) &= p(c) \prod_{t=1}^{T} \frac{1}{Z} p_\theta(x_t, c|x_{1:t-1}) \\ &= \frac{1}{Z^{T-1}} \left(\prod_{t=1}^{T} p_\theta(x_t, c|x_{1:t-1}) \right) \propto \prod_{t=1}^{T} p_\theta(x_t, c|x_{1:t-1})\,. \end{aligned} \tag{6}$$

Here, we also assume that Z^{T-1} is a constant. This is true due to the *constant time assumption*, which states that T stays the same for all generation tasks. Following Eq. 2, we can similarly decompose and represent $p_\theta(x_t, c|x_{1:t-1})$ on the right-hand side in Eq. 6 in an energy function form

$$p_\theta(\boldsymbol{x}, c) \propto \prod_{t=1}^{T} p_\theta(x_t, c|x_{1:t-1}) \propto \prod_{t=1}^{T} \underbrace{p_\theta(x_t|x_{1:t-1}) e^{-E_\theta(c|x_{1:t})}}_{P_\theta(x_t, c|\cdot)}\,. \tag{7}$$

We denote $P_\theta(x_t, c|\cdot)$ as th *step-wise joint EBM* and Eq. 7 is the step-wise joint EBM approach to CTG. The key difference between 2 and 7 is that instead of building a joint EBM with the entire sequence, our joint EBM is formulated in each time step. This formulation enables very efficient sampling: instead of sampling an entire sequence $\boldsymbol{x}$ to satisfy the control c, which is slow, we now sample x_t at each time step, which is easy because the step-wise joint EBM can be *normalized* and computed *exactly*:

$$p_\theta(x_t, c|\cdot) = \frac{1}{Z_\theta} P_\theta(x_t, c|\cdot), \quad \text{where } Z_\theta = \sum_{x_t \in \mathcal{V}} P_\theta(x_t, c|\cdot)\,, \tag{8}$$

where $Z_\theta = \sum_{x_t \in \mathcal{V}} P_\theta(x_t, c|\cdot)$ can be computed by enumerating all possible x_t, which is a finite set. Having an normalized energy function also removes the need to perform sampling via Langevin Dynamics (Welling & Teh, 2011) which tend to be slow and unstable. To further improve computational efficiency, we can approximate Z_θ in Eq. 8 by using only the top K x_t's with the highest probabilities, an idea similar to top-k sampling (Fan et al., 2018; Radford et al., 2019):

$$\widetilde{p}_\theta(x_t, c|\cdot) = \frac{1}{Z_k} P_\theta(x_t, c|\cdot) \text{ for } x_t \in \mathcal{S}_k := \left\{ x_t | \underset{k=1:K}{\text{argmax}}\ P_\theta(x_t^{(k)}, c|\cdot) \right\} \tag{9}$$

where $Z_k = \sum_{x_t \in \mathcal{S}_k} P_\theta(x_t, c|\cdot)$.

Remarks. The independence and constant time assumptions enable us to approximate the computationally expensive sequence-level EBM methods in Eq. 2 as our efficient step-wise EBM framework in Eq. 7. The trade-off we have made here is that our step-wise EBM framework models a strictly smaller class of functions and is therefore less expressive than sequence-level EBM methods. In particular, the independence assumption is a strong assumption, requiring $p_\theta(c|x_{1:t-1}) = Z\ \forall t$. The constant time assumption is milder and more natural: while allowing varying T enables more generation flexibility, we observe that 1) many previous works on CTG perform and evaluate generation using a fixed T (Khalifa et al., 2021; Ziegler et al., 2019; Liu et al., 2021; Dathathri et al., 2020) and 2) in practice, different generation lengths can be achieved simply by selecting a portion of the length T generation for shorter texts or concatenating multiple generations of length T for longer texts. However, empirically, we observe that various methods under our step-wise EBM framework, including a number of existing CTG approaches, perform surprisingly well, with performance rivaling sequence-level EBM methods while saving long hours of training/fine-tuning required for sequence-level EBM methods. We thus argue that, despite our strong assumptions, our step-wise EBM framework is an efficient and effective approximation to sequence-level EBMs. Further analysis on the small performance gap between the two is left for future work.

Generality of our framework. We have now obtained a step-wise EBM framework for CTG. Our framework simplifies the CTG problem into crafting suitable energy functions $E_\theta(c|x_{1:t})$ and offers significant flexibility in developing new CTG methods in the form of new energy function forms. Indeed, through the lens of energy functions, our framework is highly general and unifies many existing step-wise CTG methods by viewing them as specific energy function designs. For example, we can recover the control method in Dathathri et al. (2020) by implementing our energy function as an latent classifier that takes the LM's hidden states as input and running the approximated Langevin dynamics on the resulting step-wise EBM. We can recover the methods proposed in Krause et al. (2020) and Liu et al. (2021) by taking our energy function as an attribute classifier computed by two other attribute conditional LMs. We can also recover the method proposed in Yang & Klein (2021) by instantiating the energy function as a weighted combination of multiple attribute classifiers. Detailed derivations and the recovery of a few other methods (Xu et al., 2021; Lin & Riedl, 2021) under our framework are available in Supplementary Material A. We additionally include more extensive review of related literature in Supplementary Material B.

2.3 Our Energy Function Design

We exploit the idea of "control with keywords" that strategically uses the selected keywords (tokens) to steer generation towards the desired attribute. Prior research (Ribeiro et al., 2016; Wallace et al., 2018; 2019; Wiegreffe & Pinter, 2019; Kumar et al., 2021) has shown that certain words carry important attribute-sensitive information; their presence or absence in the generated text can thus indicate the text's attribute. In fact, Li et al. (2018) and Lu et al. (2021) have successfully applied this "control with keywords" idea for the text style transfer and lexical control tasks, respectively. These prior works provide strong empirical support for the adaption of this idea to CTG.

Suppose each attribute c is associated with a unique set of K_c keywords $\mathcal{V}_c$, i.e., $\mathrm{card}(\mathcal{V}_c) = K_c$. We then define a generic form of our energy function in Eq. 7

$$E(c|x_{1:t}) = \lambda_t \psi_c(x_{1:t}) + \gamma_t \psi_{\bar{c}}(x_{1:t}) \tag{10}$$

where λ_t and γ_t are the weighting parameters for the desired and undesired attributes, respectively. Also, ψ_c and $\psi_{\bar{c}}$ are indicator functions for the desired and undesired attributes, respectively. Intuitively, our energy function simply re-weighs probabilities of the keywords associated with the desired and undesired attribute. This simple generic form enables a highly flexible design of specific energy functions because we can easily adjust the parameters λ_t, γ_t, ψ_c, and $\psi_{\bar{c}}$ to suit different CTG scenarios. Furthermore, in contrast to existing step-wise CTG approaches (Dathathri et al., 2020; Krause et al., 2020; Liu et al., 2021; Lin & Riedl, 2021; Yang & Klein, 2021; Xu et al., 2021), our method requires no attribute classifiers or conditional LMs that increase the computational overhead, sometimes significantly. This enables our method to be significantly more efficient than the above methods. Below, we instantiate Eq. 10 for a number of practical CTG tasks.

For word and topic control, we use the following configurations:

$$\lambda_t = -t\lambda\,, \quad \psi_c(x_{1:t}) = \mathbb{1}\{x_t \in \mathcal{V}_c\} \cdot \mathbb{1}\{\mathcal{V}_c \cap x_{1:t-1} = \emptyset\}\,, \quad \gamma_t = 0\,, \quad \psi_{\bar{c}}(x_{1:t}) = 0\,. \tag{11}$$

Note that $K_c = 1$ in the word control scenario. Intuitively, we up-weigh the probability of words in $\mathcal{V}_c$ if it has not yet appeared in the previously generated tokens $x_{1:t-1}$. The energy function increases linearly in t, such that we up-weigh words in $\mathcal{V}_c$ more towards the completion of the generation if none of keywords still has not appeared. The hyper-parameter λ controls the aggressiveness of the control: the higher the value, the more aggressive we up-weigh the sampling of keywords.

For sentiment control, we use the following configurations:

$$\lambda_t = -\lambda \;\forall\; t\,, \quad \psi_c(x_{1:t}) = \mathbb{1}\{x_t \in \mathcal{V}_c\}\,, \quad \gamma_t = \infty \;\forall\; t\,, \quad \psi_{\bar{c}}(x_{1:t}) = \mathbb{1}\{x_t \in \mathcal{V}_{\bar{c}}\}\,. \tag{12}$$

Intuitively, we up-weigh the keywords of the desired sentiment with the constant λ and avoid the keywords of the undesired sentiment. To construct $\mathcal{V}_c$ and $\mathcal{V}_{\bar{c}}$, we first predict each word's sentiment and pre-compute its sentiment probability with an external sentiment classifier. We then choose two thresholds τ_c and $\tau_{\bar{c}}$ for the desired and undesired sentiment, respectively, and include words above each threshold into either $\mathcal{V}_c$ or $\mathcal{V}_{\bar{c}}$. These hyper-parameters $\lambda, \tau_c, \tau_{\bar{c}}$ together control the aggressiveness of satisfying the desired sentiment.

For detoxification, we denote $\bar{c}$ as the toxic attribute and use the following configurations:

$$\lambda_t = 0 \;\forall\; t\,, \quad \psi_c(x_{1:t}) = 0\,, \quad \gamma_t = \infty \;\forall\; t\,, \quad \psi_{\bar{c}}(x_{1:t}) = \mathbb{1}\{x_t \in \mathcal{V}_{\bar{c}}\}\,. \tag{13}$$

Intuitively, this configuration avoids words that are toxic. Similar to sentiment control, we construct $\mathcal{V}_{\bar{c}}$ by pre-computing each word's toxicity with an external toxicity classifier and then selecting words above a preset toxicity score threshold $\tau_{\bar{c}}$ into $\mathcal{V}_{\bar{c}}$.

For composition of multiple controls, i.e., imposing multiple controls simultaneously, we extend the compositionality property in Eq. 3 and construct the following energy function:

$$E(c_1, \ldots, c_N|x_{1:t}) = \sum_{i=1}^{N} E_i(c_i|x_{1:t})\,, \tag{14}$$

which enables a simple summation of multiple energy functions $E_i(c_i|x_{1:t})$ for $i = 1, \cdots, N$ such as those instantiated above. Note that each energy function $E_i(c_i|x_{1:t})$ in 14 may correspond to a separate control scenario, i.e., a composition of topic control and sentiment control.

3 Experiments

We now present extensive evaluation results to validate our proposed method, covering a variety of control scenarios: i) single word control, ii) topic control, iii) sentiment control, iv) detoxification, and v) composition of multiple controls. The overall theme of the experiments is threefold. First, our method achieves competitive (sometimes superior) performance on each evaluation metric compared to baselines (Sections 3.1 and 3.2). Second, our method nicely trades off control satisfaction and generation quality, which can be easily user-adjusted by changing the hyper-parameter in the energy functions (Sections 3.1 – 3.4). Third, our method uses (sometimes significantly) less computational resources compared to baselines (Section 3.5). Supplementary Materials C and D contain additional results including qualitative evaluations, analyses, and other control settings.

Table 2: Word control and topic control experiment results. Our method guarantees that the single control word or any keyword in a topic appears in the generation, and also achieves high language quality and diversity.

Target Control	Method	Ctrl. Satisfaction (↑) Word/Topic Acc.	Fluency (↓) perplexity	diversity (↑) dist-1	dist-2	dist-3
Word Control	REINFORCE (Williams, 1992)	**1.000**	420.09	0.268	0.476	0.612
	REINFORCE_Px (Khalifa et al., 2021)	0.357	19.66	0.727	0.798	0.790
	PPO (Ziegler et al., 2019)	0.978	11.22	0.811	0.942	**0.936**
	GDC (Khalifa et al., 2021)	0.931	**10.97**	0.811	0.936	0.930
	PPLM (Dathathri et al., 2020)	0.212	11.28	0.810	0.933	0.927
	FUDGE (Yang & Klein, 2021)	0.480	12.29	0.730	0.854	0.858
	ours	**1.000**	12.45	**0.822**	**0.944**	**0.936**
Topic Control	REINFORCE (Williams, 1992)	**1.000**	420.086	0.268	0.476	0.612
	REINFORCE_Px (Khalifa et al., 2021)	0.338	20.351	0.822	0.898	0.882
	PPO (Ziegler et al., 2019)	0.967	**9.720**	0.812	0.935	0.930
	GDC (Khalifa et al., 2021)	0.830	10.075	0.805	0.934	0.930
	PPLM (Dathathri et al., 2020)	0.400	10.577	0.812	0.937	0.931
	CTRL (Shirish Keskar et al., 2019)	0.392	48.693	**0.930**	**0.970**	**0.946**
	FUDGE (Yang & Klein, 2021)	0.686	12.900	0.797	0.926	0.922
	ours	**1.000**	10.928	0.807	0.937	0.932

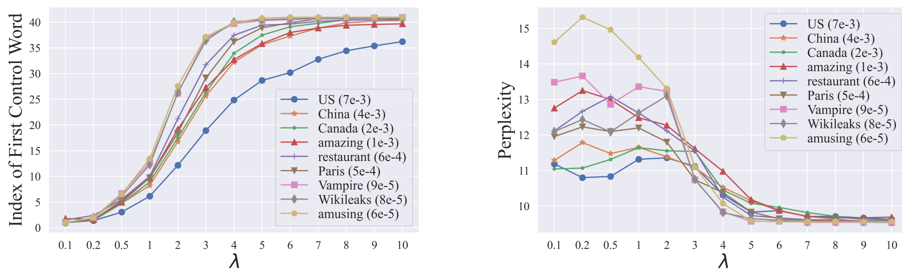

Figure 1: Analysis of the hyper-parameter λ of Eq. 11 in word control. **Left**: how λ influences control satisfaction, i.e., how early the control word appears in the generated text and whether it appears at all. **Right**: how λ influences and language quality, i.e., perplexity. In general, we observe a trade-off between control satisfaction and language quality.

Baselines and setup. We consider three categories of baselines. The first category, including **REINFORCE** (Williams, 1992), **REINFORCE_Px** (Khalifa et al., 2021), **PPO** (Ziegler et al., 2019), and **GDC** (Khalifa et al., 2021), performs control under the sequence-level EBM framework by fine-tuning an LM through reinforcement learning techniques. The second category, including **CTRL** (Shirish Keskar et al., 2019) and **DAPT** (Gururangan et al., 2020), adopts the "prompt engineering" approach (Liu et al., 2021) and fine-tunes an LM using pre-specified prompts as the control signal. The third category, including **PPLM** (Dathathri et al., 2020), **GeDi** (Krause et al., 2020), **FUDGE** (Yang & Klein, 2021), **PnB** (Lin & Riedl, 2021), and **DExperts** (Liu et al., 2021), performs step-wise control under the guidance of attribute classifiers during generation. Throughout the experiments, whenever possible, we use a small GPT-2 as the base LM, use the same top-k sampling where $k = 10$, and disable n-gram penalty (Paulus et al., 2018; Klein et al., 2017) for all methods. Note that not all baselines are suitable for all five control scenarios and we thus use a subset of baselines in each control scenario. More details are available in Supplementary Material C.

Metrics. We consider three types of metrics to evaluate i) control satisfaction, ii) language quality, iii) generation diversity. For control satisfaction, we use the **accuracy** metric. Specifically, in single word and topic controls, accuracy denotes the percentage of generated samples that contain the control word and any of the keywords associated with a topic, respectively. In sentiment control, accuracy is the percentage of sentiment predictions, computed by external sentiment (Wolf et al., 2020), that matches with the desired control sentiment. In detoxification, accuracy is the averaged probabilities of toxicity of the generated samples, computed using an external toxicity classifiers (Perspective, 2021). For language quality, we use the average **perplexity** over all generated samples, computed using a pre-trained small GPT-2. For generation diversity, we use the **distinct-N** metric (Li et al., 2016) with $N \in \{1, 2, 3\}$, which computes the total number of distinct N-grams in each generated text, scaled by the total number of tokens and averaged over all generations. We additionally perform **human evaluations** to evaluate on control satisfaction and language quality for the sentiment control and detoxification experiments where the control is subtle and automatic evaluations may be insufficient. We defer the detailed setup to Supplementary Material C.3.

3.1 Single Word and Topic Controls

In word control, we follow Khalifa et al. (2021) and choose nine control words with varying rarity levels, including "US", "China", "Canada", "restaurant", "Paris", "Vampire", "Wikileaks", and

Table 3: Sentiment control experiment results. Our method achieves highly competitive results, with the second best control satisfaction and the third highest language quality for both positive and negative sentiments.

Target Sentiment	Method	Ctrl. Satisfaction (↑) Sentiment Acc.	Fluency (↓) perplexity	Language Diversity (↑) dist-1	dist-2	dist-3	Human Evaluation Sentiment	Fluency (↑)
Positive	REINFORCE (Williams, 1992)	**1.000**	19.071	0.213	0.304	0.367	-	-
	REINFORCE_Px (Khalifa et al., 2021)	0.533	38.118	0.701	0.848	0.872	-	-
	PPO (Ziegler et al., 2019)	0.998	13.173	0.748	0.921	0.931	**3.273**	1.573
	GDC (Khalifa et al., 2021)	0.883	**11.742**	0.757	0.915	0.924	2.863	1.653
	DExperts (Liu et al., 2021)	0.977	24.244	0.766	0.888	0.886	3.047	1.657
	DAPT (Gururangan et al., 2020)	0.872	60.294	0.860	0.957	0.940	2.867	1.610
	GeDi (Krause et al., 2020)	0.845	102.157	0.859	0.958	0.936	2.500	1.133
	CTRL (Shirish Keskar et al., 2019)	0.810	47.207	**0.922**	**0.974**	**0.953**	3.183	**1.750**
	PPLM (Dathathri et al., 2020)	0.815	26.694	0.705	0.878	0.892	2.747	1.303
	ours	0.932	17.866	0.729	0.881	0.897	2.877	1.413
	GPT-2	0.570 (positive)	11.189	0.764	0.920	0.923	-	-
Negative	REINFORCE (Williams, 1992)	**0.997**	72.789	0.614	0.803	0.850	-	-
	REINFORCE_Px (Khalifa et al., 2021)	0.790	19.513	0.587	0.761	0.811	-	-
	PPO (Ziegler et al., 2019)	0.917	12.823	0.735	0.912	0.927	1.147	1.383
	GDC (Khalifa et al., 2021)	0.705	**11.347**	0.759	0.912	0.919	1.680	1.497
	DExperts (Liu et al., 2021)	0.992	29.635	0.817	0.949	0.937	**0.913**	1.053
	DAPT (Gururangan et al., 2020)	0.880	62.0901	0.875	0.957	0.938	1.433	1.480
	GeDi (Krause et al., 2020)	0.965	66.571	0.890	0.958	0.938	1.280	1.143
	CTRL (Shirish Keskar et al., 2019)	0.845	48.057	**0.921**	**0.974**	**0.953**	1.190	**1.693**
	PPLM (Dathathri et al., 2020)	0.635	21.784	0.717	0.894	0.908	1.910	1.287
	ours	0.965	16.665	0.712	0.883	0.901	1.237	1.317

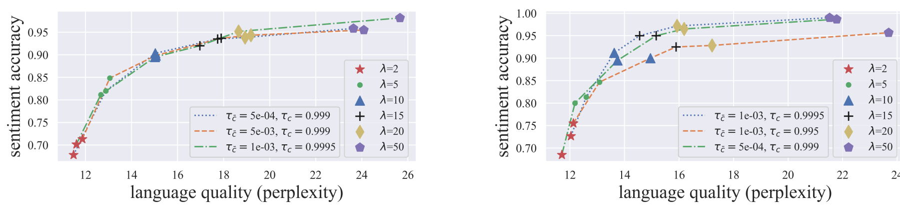

Figure 2: Analysis of how the hyper-parameters λ, τ_c, $\tau_{\bar{c}}$ in Eq. 12 influence control satisfaction and language quality for positive (**left** plot) and negative (**right** plot) sentiments. Larger λ, τ_c and lower $\tau_{\bar{c}}$ generally result in better control satisfaction but worse language quality.

"amusing". We list the above control words in the order from the most common to the rarest. In topic control, we follow Dathathri et al. (2020) and use nine topics including "science", "fantasy", "space", "politics", "military", "religion", "computers", "kitchen", and "legal". We use the same topic keywords as in (Dathathri et al., 2020). Each method generates 100 samples of 40 tokens from the empty prompt "`|<endoftext>|`", totalling 900 generations for each control setting.

Results. Table 2 reports the evaluation results averaged over all 900 generations for each method in word control and topic control, respectively. In terms of control satisfaction, we see that only REINFORCE and our method fully satisfy the control requirement. However, REINFORCE tends to produce degenerate samples consisting of only the control word, as evidenced by its high perplexity and low dist-N scores. In terms of language quality and diversity, our method generates samples of quality on par with the best performing baselines. Some baselines including PPLM and FUDGE achieve high language quality and diversity but fail to satisfy the control attribute. These results show that our method can generate high-quality samples and satisfy the desired control simultaneously whereas most baselines cannot.

Analysis of hyper-parameters. We perform an empirical analysis to better understand how the hyper-parameter λ in Eq. 11 impacts generated samples. Recall that λ controls the aggressiveness of the control so we would expect that λ impacts i) how early the control word first appears in the generation and ii) the language quality of the generation. By varying λ from 0.1 to 10, we plot in Figure 1 the averaged first index of the control word in the generated samples and the averaged perplexity of the generated samples. Overall, we observe a trade-off between control satisfaction and language quality: the more aggressive (larger λ) we wish to impose the control, the earlier the control word appears in the generation, and the worse the language quality. When λ is small, the perplexity approaches that of GPT-2 but the word may not even appear in the generation, indicating a failure to satisfy the control. Another observation is that for more frequent words, we need a smaller λ to achieve high control satisfaction and language quality at the same time. These insights provide a useful guideline for adjusting λ to achieve the desired trade-off in different word control scenarios. We include similar analyses for topic controls in Supplementary Material D.

3.2 Sentiment Control

We follow Khalifa et al. (2021) and generate texts with two sentiments, i.e., positive and negative, from six prompts including "The horse", "The pizza", "The lake", "The chicken", "The potato",

Table 4: Language detoxification results evaluated on samples generated from toxic prompts. Our method achieves the second lowest text toxicity, third best perplexity, and competitive diversity scores, indicating it best balances the effective detoxification and high-quality language generation.

Method	Toxicity (↓)	Perplexity (↓)	diversity (↑)			Human Evaluation	
			dist-1	dist-2	dist-3	Toxicity (↓)	Fluency (↑)
GPT-2 (Radford et al., 2019)	0.598	14.777	0.730	0.876	0.889	-	-
REINFORCE (Williams, 1992)	0.239	20.095	0.631	0.829	0.874	-	-
REINFORCE_Px (Khalifa et al., 2021)	0.469	17.702	0.630	0.780	0.820	-	-
PPO (Ziegler et al., 2019)	0.274	**13.893**	0.776	0.923	0.927	1.123	**1.650**
GDC (Khalifa et al., 2021)	0.523	14.520	0.734	0.878	0.888	1.237	1.523
DExperts (Liu et al., 2021)	0.229	25.244	0.746	0.839	0.832	1.090	1.433
DAPT (Gururangan et al., 2020)	0.277	95.319	**0.869**	**0.954**	**0.933**	1.133	1.497
GeDi (Krause et al., 2020)	**0.080**	389.582	0.669	0.783	0.801	**1.030**	0.283
PPLM (Dathathri et al., 2020)	0.503	21.752	0.773	0.906	0.909	1.293	1.633
ours	0.206	14.867	0.740	0.889	0.900	1.150	1.608
GPT-3	0.547	31.302	**0.776**	**0.891**	**0.884**	-	-
GPT-3 + DExperts	0.346	**12.368**	0.650	0.811	0.845	1.197	**1.603**
GPT-3 + ours	**0.261**	18.976	0.611	0.766	0.807	**1.133**	1.590

and the empty prompt. Each method generates 100 samples of 40 tokens from one of the above six prompts, totalling 600 generations for each sentiment.

Results. Table 3 reports evaluation results averaged over all 600 generations for each method. Our method achieves highly competitive results, with the second best control satisfaction and the third highest language quality for both positive and negative sentiment controls. Some baselines (Dexperts, DAPT, GeDi, CTRL) are able to achieve better language quality by using more resources than our method during generation. For example, CTRL and DAPT use a base LM (GPT-2 Large) with 6.6 times the parameters than the one we use (GPT-2 small). GeDi and Dexperts and GeDi use additional conditional generative LMs that are even larger than the base LM. The generated samples of our method also have a decent diversity given that we do not use any extra resources other than the base LM and pre-constructed sentiment keywords. Human evaluations also suggest that our method achieves competitive performance in terms of sentiment satisfaction and fluency.

Analysis of hyper-parameters. We perform an empirical analysis to better understand the effects of various hyper-paremters ($\lambda, \tau_c, \tau_{\bar{c}}$) in Eq. 12 on the trade-off between control satisfaction and language quality. Figure 2 demonstrates this trade-off with varying $\lambda, \tau_c, \tau_{\bar{c}}$ values for both positive and negative sentiments. We find that, the more aggressive we set these hyper-parameters, i.e., higher λ, τ_c and lower $\tau_{\bar{c}}$, the better control satisfaction but worse language quality. This observation confirms our intuition: more aggressive values of these hyper-parameters increase the chance that the words indicative of the desired sentiment appear in the generated text but may also result in their repetitive appearance. In the extreme case, the generation may completely neglect other words that have no desired sentiment, leading to worse language quality and diversity. In practice, $\lambda, \tau_c, \tau_{\bar{c}}$ serve as knobs that enables practitioners to tune the trade-off to suit different applications.

3.3 DETOXIFICATION

We generate 10 samples from 100 toxic prompts in the RealToxicPrompt dataset (Gehman et al., 2020) for each method and show the detodification results in Table 4. Our approach achieves the second lowest text toxicity, third best perplexity, and competitive diversity scores. The method GeDi that achieves the lowest text toxicity score has a rather high perplexity, indicating unsatisfactory language quality. In general, our approach best balances the effective detoxification and high-quality language generation. We additionally test our method's ability in detoxification for OpenAI's GPT-3 model "ada", and achieves better performance compared to DExperts-guided GPT-3. Human evaluations further corroborate the above observations, where our method achieves third lowest toxicity and third best fluency as determined by human evaluators when using GPT-2 as base LM and comparable performance with DExperts when using GPT-3 as base LM. Finally, we evaluate the generated samples with 100 nontoxic prompts and the results resemble the observations in Table 4. We also analyze the effect of the hyper-parameter $\tau_{\bar{c}}$ and observe that, while higher value correlates with more text toxicity, it does not significantly impact perplexity. This observation suggests that removing toxic words does not impair language quality and that our keyword-based energy function is highly effective in the detoxification task. Results on nontoxic prompts and hyper-parameter analysis are in Supplementary Material D.

3.4 COMPOSITIONS OF MULTIPLE CONTROLS

We consider three compositional control settings, including i) combinations of two topics, ii) two topics and positive sentiment, and ii) two topics and negative sentiment. For each of the three

Table 5: Compositional control experiment results. Our method achieves the best balance of good control satisfaction and high language quality and diversity.

Controls	Method	Ctrl. Satisfaction (↑) topic acc.	Ctrl. Satisfaction (↑) sentiment acc.	Fluency (↓) perplexity	Diversity (↑) dist-1	Diversity (↑) dist-2	Diversity (↑) dist-3
two topics	PnB (Lin & Riedl, 2021)	0.287	-	75.934	**0.842**	**0.967**	**0.958**
	PPLM (Dathathri et al., 2020)	0.171	-	9.927	0.732	0.926	0.942
	GDC (Khalifa et al., 2021)	0.621	-	11.569	0.803	0.935	0.931
	ours	**0.988**	-	**9.837**	0.744	0.930	0.943
two topics & pos. sentiment	PPLM (Dathathri et al., 2020)	**1**	0.605	23.222	0.268	0.366	0.464
	GDC (Khalifa et al., 2021)	0.540	0.784	14.274	**0.793**	**0.931**	**0.929**
	ours	**0.949**	**0.845**	**12.982**	0.699	0.880	0.908
two topics & neg. sentiment	PPLM (Dathathri et al., 2020)	**1**	0.677	24.688	0.274	0.371	0.470
	GDC (Khalifa et al., 2021)	0.486	0.745	15.060	**0.766**	0.916	0.920
	ours	0.985	**0.846**	**11.529**	0.733	**0.923**	**0.940**

Table 6: Computational efficiency comparing several baselines to ours with different base LM sizes where "S" and "L" denote small and large GPT-2 models, respectively. Our method requires no training/fine-tuning, uses the least model parameters, and achieves fastest generation among all baselines, regardless of base LM sizes.

Metrics	PPO (S)	GDC (S)	GeDi (S)	**Ours (S)**	CTRL (L)	PPLM (L)	DExperts (L)	**Ours (L)**
Training Time (hrs) (↓)	9+	100+	0	0	0	0	**0**	**0**
Generation Speed (tokens/sec) (↑)	83.05	81.35	33.19	**97.74**	1.95	0.73	11.89	**39.59**
Model Size (millions of params.) (↓)	124.44	124.44	479.26	**124.44**	1637.96	774.03	2322.09	**774.03**

settings, we consider three sets of two topics that have no overlapping topic keywords, including {space, military}, {politics, religion}, and {fantasy, computers}. Each method generates 500 samples of 60 tokens from the empty prompt for each set of topics, totalling 1500 generations for each setting. Table 5 reports the results, averaged over all 1500 generations for each compositional control setting. Our method better satisfies both topic and sentiment controls simultaneously and achieves competitive or superior generation quality and diversity compared to baselines. In contrast, the baselines either do not satisfactorily fulfill all controls (PnB, GDC) or suffer from language quality and diversity (PPLM). Therefore, our method achieves the best balance of good control satisfaction and high language quality and diversity, an observation that corroborates earlier empirical results.

3.5 Computational Efficiency

Finally, we demonstrate the superior computational efficiency of our approach to supplement its outstanding CTG capability. We compute the model size (number of parameters), training/fine-tuning time (hours), and generation speed (in seconds per token) as the efficiency metrics. In Table 6, we compare our approach with the top performing baselines. We can see that our approach requires no training/fine-tuning, uses the least number of model parameters, and generates text the fastest among all baselines. This is especially significant considering the strong baselines we compare with can hardly maintain both the good CTG performance and high computational efficiency. For example, GDC also nicely balances all metrics for most of the CTG experiments but requires 100+ hours of fine-tuning. CTRL often achieves higher language diversity but use 10x more model parameters than ours. DExperts achieves slightly better performance in the sentiment control experiment but requires 20x more model parameters and are 8x slower in generation than ours. All the above evaluation results suggest that our approach is highly appealing as an efficient and effective plug-and-play method for a wide range of practical CTG scenarios.

4 Conclusions

In this work, we have studied the problem of CTG and proposed a step-wise EBM framework for tackling this problem. Our general framework not only unifies a number of existing CTG approaches but also inspires an efficient, effective, and flexible energy function design based on keywords associated with a specific control attribute. We have demonstrated the wide applicability of our proposed control approach on a five different CTG tasks. Our approach consistently achieves competitive performance, provides an tunable hyper-parameter to easily adjust the trade-off between control satisfaction and language quality, and adds an infinitesimal amount of computational overhead.

Our work suggests that a range of practical CTG tasks can be tackled with methods as simple and efficient as ours. However, for CTG tasks with more nuanced constraints, the "control with keywords" intuition that underlies our energy function design could fail. A future research avenue is to investigate more elaborate energy function designs that will enable more complex CTG tasks, such as text style transfer (Han & Lundin, 2021) and text rewriting (Coenen et al., 2021).

REFERENCES

Anton Bakhtin, Yuntian Deng, Sam Gross, Myle Ott, Marc'Aurelio Ranzato, and Arthur Szlam. Residual energy-based models for text. *J. Mach. Learn. Res.*, 22(40):1–41, 2021.

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proc. ACM Conf. Fairness Accountability Transparency*, pp. 610–623, 2021.

Sumanta Bhattacharyya, Amirmohammad Rooshenas, Subhajit Naskar, Simeng Sun, Mohit Iyyer, and Andrew McCallum. Energy-based reranking: Improving neural machine translation using energy-based models. In *Proc. Annu. Meeting Assoc. Comput. Linguistics and Int. Joint Conf. Natural Lang. Process.*, pp. 4528–4537, August 2021.

Shikha Bordia and Samuel R. Bowman. Identifying and reducing gender bias in word-level language models. In *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Student Res. Workshop*, pp. 7–15, June 2019.

Tom Brown et al. Language Models are Few-Shot Learners. *arXiv e-prints*, art. 2005.14165, May 2020.

Shuyang Cao and Lu Wang. Inference time style control for summarization. In *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Human Lang. Technol.*, pp. 5942–5953, June 2021.

Alvin Chan, Ali Madani, Ben Krause, and Nikhil Naik. Deep Extrapolation for Attribute-Enhanced Generation. *arXiv e-prints*, art. 2107.02968, July 2021.

Alvin Chan, Yew-Soon Ong, Bill Pung, Aston Zhang, and Jie Fu. Cocon: A self-supervised approach for controlled text generation. In *Proc. Int. Conf. Learn. Representations*, 2021.

Tong Che, Ruixiang ZHANG, Jascha Sohl-Dickstein, Hugo Larochelle, Liam Paull, Yuan Cao, and Yoshua Bengio. Your gan is secretly an energy-based model and you should use discriminator driven latent sampling. In *Proc. Neural Inf. Process. Syst.*, volume 33, pp. 12275–12287, 2020.

Andy Coenen, Luke Davis, Daphne Ippolito, Emily Reif, and Ann Yuan. Wordcraft: a Human-AI Collaborative Editor for Story Writing. *arXiv e-prints*, art. 2107.07430, July 2021.

Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. Plug and play language models: A simple approach to controlled text generation. In *Proc. Int. Conf. Learn. Representations*, 2020.

Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. *arXiv e-prints*, art. 1805.04833, May 2018.

Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. RealToxicityPrompts: Evaluating neural toxic degeneration in language models. In *Proc. Conf. Empirical Methods Natural Lang. Process.*, pp. 3356–3369, November 2020.

Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. Hafez: an interactive poetry generation system. In *Proc. Annu. Meeting Assoc. Comput. Linguistics*, pp. 43–48, July 2017.

Will Grathwohl, Kuan-Chieh Wang, Joern-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. In *Proc. Int. Conf. Learn. Representations*, 2020.

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. Don't stop pretraining: Adapt language models to domains and tasks. In *Proc. Annu. Meeting Assoc. Comput. Linguistics*, pp. 8342–8360, July 2020.

Xing Han and Jessica Lundin. Multi-pair text style transfer for unbalanced data via task-adaptive meta-learning. In *Proc. Workshop Meta Learn. and Its Appl. Natural Lang. Process.*, pp. 28–35, August 2021.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *Proc. Int. Conf. Learn. Representations*, 2020.

Muhammad Khalifa, Hady Elsahar, and Marc Dymetman. A distributional approach to controlled text generation. In *Proc. Int. Conf. Learn. Representations*, 2021.

Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. OpenNMT: Open-Source Toolkit for Neural Machine Translation. *arXiv e-prints*, art. 1701.02810, January 2017.

Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. GeDi: Generative Discriminator Guided Sequence Generation. *arXiv e-prints*, art. 2009.06367, September 2020.

Sachin Kumar, Eric Malmi, Aliaksei Severyn, and Yulia Tsvetkov. Controlled Text Generation as Continuous Optimization with Multiple Constraints. *arXiv e-prints*, art. 2108.01850, August 2021.

Sawan Kumar, Kalpit Dixit, and Kashif Shah. Interpreting text classifiers by learning context-sensitive influence of words. In *Proc. Workshop Trustworthy Natural Lang. Process.*, pp. 55–67, June 2021.

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A diversity-promoting objective function for neural conversation models. In *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Human Lang. Technol.*, pp. 110–119, June 2016.

Juncen Li, Robin Jia, He He, and Percy Liang. Delete, retrieve, generate: a simple approach to sentiment and style transfer. In *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Human Lang. Technol.*, pp. 1865–1874, June 2018.

Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. Towards understanding and mitigating social biases in language models. In *Proc. Int. Conf. Mach. Learn.*, volume 139, pp. 6565–6576, 18–24 Jul 2021.

Zhiyu Lin and Mark Riedl. Plug-and-blend: A framework for controllable story generation with blended control codes. In *Proc. Workshop Narrative Understanding*, pp. 62–71, June 2021.

Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. DExperts: Decoding-time controlled text generation with experts and anti-experts. In *Proc. Annu. Meeting Assoc. Comput. Linguistics and Int. Joint Conf. Natural Lang. Process.*, pp. 6691–6706, August 2021.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *arXiv e-prints*, art. 2107.13586, July 2021.

Ruibo Liu, Jason Wei, Chenyan Jia, and Soroush Vosoughi. Modulating language models with emotions. In *Proc. Findings Assoc. Comput. Linguistics*, pp. 4332–4339, August 2021.

Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. NeuroLogic decoding: (un)supervised neural text generation with predicate logic constraints. In *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Human Lang. Technol.*, pp. 4288–4299, Online, June 2021.

Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Proc. Annu. Conf. Int. Speech Commun. Assoc.*, volume 2010, pp. 1045–1048, 2010.

Tomáš Mikolov, Stefan Kombrink, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Extensions of recurrent neural network language model. In *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, pp. 5528–5531, 2011.

Romain Paulus, Caiming Xiong, and Richard Socher. A deep reinforced model for abstractive summarization. In *Proc. Int. Conf. Learn. Representations*, 2018.

Nanyun Peng, Marjan Ghazvininejad, Jonathan May, and Kevin Knight. Towards controllable story generation. In *Proc. Workshop Storytelling*, pp. 43–49, June 2018.

Perspective. Using machine learning to reduce toxicity online, 2021. URL `https://www.perspectiveapi.com/`.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should i trust you?": Explaining the predictions of any classifier. In *Proc. ACM SIGKDD Int. Conf Knowl. Discovery Data Mining*, pp. 1135–1144, 2016.

Alexis Ross, Tongshuang Wu, Hao Peng, Matthew E. Peters, and Matt Gardner. Tailor: Generating and Perturbing Text with Semantic Controls. *arXiv e-prints*, art. 2107.07150, July 2021.

Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. The woman worked as a babysitter: On biases in language generation. In *Proc. Conf. Empirical Methods in Natural Lang. Process. and Int. Joint Conf. Natural Lang. Process.*, pp. 3407–3412, November 2019.

Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. CTRL: A Conditional Transformer Language Model for Controllable Generation. *arXiv e-prints*, art. 1909.05858, September 2019.

Eric Wallace, Shi Feng, and Jordan Boyd-Graber. Interpreting neural networks with nearest neighbors. In *Proc. EMNLP Workshop BlackboxNLP Analyzing and Interpreting Neural Networks for NLP*, pp. 136–144, November 2018.

Eric Wallace, Jens Tuyls, Junlin Wang, Sanjay Subramanian, Matt Gardner, and Sameer Singh. AllenNLP interpret: A framework for explaining predictions of NLP models. In *Proc. Conf. Empirical Methods in Natural Lang. Process. and Int. Joint Conf. Natural Lang. Process.*, pp. 7–12, November 2019.

Zhe Wang, Wei He, Hua Wu, Haiyang Wu, Wei Li, Haifeng Wang, and Enhong Chen. Chinese poetry generation with planning based neural network. In *Proc. Int. Conf. Comput. Linguistics*, pp. 1051–1060, December 2016.

Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proc. Int. Conf. Mach. Learn.*, pp. 681–688, 2011.

Sarah Wiegreffe and Yuval Pinter. Attention is not not explanation. In *Proc. Conf. Empirical Methods in Natural Lang. Process. and Int. Joint Conf. Natural Lang. Process.*, pp. 11–20, November 2019.

Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8(3–4):229–256, May 1992.

Thomas Wolf et al. Transformers: State-of-the-art natural language processing. In *Proc. Conf. Empirical Methods Natural Lang. Process.*, pp. 38–45, Online, October 2020.

Chen Xu, Jianyu Zhao, Rang Li, Changjian Hu, and Chuangbai Xiao. Change or not: A simple approach for plug and play language models on sentiment control. *Proc. AAAI Conf. Artif. Intell.*, 35(18):15935–15936, May 2021.

Kevin Yang and Dan Klein. FUDGE: Controlled text generation with future discriminators. In *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Human Lang. Technol.*, pp. 3511–3535, June 2021.

Hongyu Zang and Xiaojun Wan. Towards automatic generation of product reviews from aspect-sentiment scores. In *Proc. Int. Conf. Natural Lang. Gener.*, pp. 168–177, September 2017.

Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-Tuning Language Models from Human Preferences. *arXiv e-prints*, art. 1909.08593, September 2019.

Supplementary Material

A How our Step-wise EBM framework Unifies Some Prior Works

In this section, we derive in detail how some prior works are special instances of our step-wise EBM framework. We reminder the reader that, throughout our paper, we use $p_\theta(\cdot)$ to denote proper probability distribution and $P_\theta(\cdot)$ to denote unnormalized distribution.

A.1 Relation to Dathathri et al. (2020)

To obtain the sampling method in Dathathri et al. (2020), we need to make a few approximations and invoke the reparametrization trick. Let $o_t = g(H_{t-1}, x_{t-1})$ where g is the language model that defines $p_\theta(\boldsymbol{x})$, H_{t-1} is the history (key, value) pairs at t, and o_t is used to compute the multinomial distribution parameter from which x_t is sampled:

$$x_t \sim \mathrm{softmax}(W^\top o_t) = f(o_t) \tag{15}$$

Further, let e_t be the unique embedding associated with x_t where $e_t = E^\top x_t$ and E is the embedding matrix. So we can also write $p_\theta(x)$ as $p_\theta(e_t)$.

First, we approximate the hard sampling x_t (or e_t) from a multinomial distribution parametrized by o_t as a soft sampling, i.e.,

$$e_t \approx E^\top f(o_t) = E^\top f(g(H_{t-1}, e_{t-1})) = E^\top \bar{g}(H_{t-1}, e_{t-1}) = \hat{e}_t\,. \tag{16}$$

Next, we define the prior distribution over H as deterministic, i.e., $p(H = H_t) = 1$ and therefore $\log p(H = H_t) = 0$. Third, we define the energy function as follows as:

$$E_\theta(c|x_{1:t}) := \lambda E_\theta(c|x_{1:t-1}, x_t, x_{t+1}) = \lambda E_\theta(c|o_{1:t-1}, o_t, o_{t+1})\,, \tag{17}$$

which depends on the current generation and 1-step soft rollout. The rollout o_{t+1} is also computed by using the soft sampling approximation in Eq. 16. Finally, we invoke the reparametrization trick, which states that sampling x_t from the distribution defined by $P(x_t, c|\cdot)$ is the same as first "sampling" H_{t-1} and then compute $\hat{e}_t$ by Eq. 16. Therefore, we have the joint EBM of H_{t-1} and c as

$$P(H_{t-1}, c) = e^{-\lambda E_\theta(c|o_{1:t-1}, o_t, o_{t+1})} \tag{18}$$

where $o_t = g(H_{t-1}, e_{t-1})$ and $o_{t+1} = E^\top \bar{g}(H_t, \hat{e}_t)$. Let $E_\theta = \log p(c|H_{t-1})$ be a classifier, we recover the update rule on H_{t-1} in Dathathri et al. (2020) by running a Langevin Dynamics (LD) on H_{t-1} with the initial H_{t-1} computed using the previous history H_{t-2} and previous token embedding e_{t-2}. Thus we see that the per-time-step sampling by updating the history proposed in Dathathri et al. (2020) corresponds to "latent"-space LD under our step-wise EBM framework with a few approximations.

The method proposed in Xu et al. (2021) takes the same idea of gradient-based steering and extends (Dathathri et al., 2020); the method therein can be similarly derived as above.

A.2 Relation to Krause et al. (2020)

Krause et al. (2020) models CTG as a joint distribution defined as follows:

$$P(x_t, c|\cdot) = p_\theta(x_t|\cdot) p(c|x_{1:t-1}, x_t)^\omega\,,$$

where ω is a weight hyper-parameter and $p(c|x_{1:t-1}, x_t)$ is an attribute classifier computed by Bayes Rule using two pre-trained attribute-conditional LMs:

$$p(c|x_{1:t}) = \frac{p(c)\prod_{j=1}^{t} p(x_j|x_{1:j}, c)}{\sum_c p(c)\prod_{j=1}^{t} p(x_j|x_{1:j}, c)}\,.$$

This method corresponds to our framework with the following energy function:

$$E_\theta(c|x_{1:t-1}) = -\omega \log p(c|x_{1:t-1}, x_t)\,.$$

A.3 Relation to Lin & Riedl (2021)

Lin & Riedl (2021) extends Krause et al. (2020) to enable compositions of N controls, with the joint distribution defined as follows:

$$P(x_t, c_1, \dots, c_N|\cdot) = p_\theta(x_t|\cdot) \prod_{i=1}^{N} p_i(c_i|x_{1:t-1}, x_t)^{\omega_i} .$$

This method corresponds to our framework with the following energy function

$$E_\theta(c_1, \dots, c_N|x_{1:t-1}) = \sum_{b=1}^{N} -\omega_i \log p_i(c_i|x_{1:t-1}, x_t) .$$

and the remaining sampling step follows from the derivation in Section A.2.

A.4 Relation to Yang & Klein (2021)

Yang & Klein (2021) proposes a control method similar to (Krause et al., 2020) by dictating the attribute classifier to depend on past and *future* generation instead of only on past generation. The method is defined as the joint distribution as follows:

$$\widetilde{P}(x_t, c|\cdot) := p_\theta(x_t|\cdot) p(c|x_{1:t+s}) ,$$

where s is the number of future time steps beyond the current time step t. This method corresponds to our framework with the following energy function:

$$E_\theta(c|x_{1:t-1}) := -\omega \log p(c|x_{1:t+s}) ,$$

where the right hand side includes additional tokens generated for s future steps.

A.5 Relation to Liu et al. (2021)

Instead of modeling the joint distribution of text and attribute, the method in Liu et al. (2021) directly augments a pre-trained base LM with a weight computed by two attribute-conditional LMs that steer the base LM towards the desired attribute and away from the undesired attribute, respectively. The method is define as follows:

$$P(x_t|x_{1:t-1}) := p_\theta(x_t|\cdot) \left(\frac{p_\theta^+(x_t|\cdot)}{p_\theta^-(x_t|\cdot)} \right)^\omega ,$$

where $p_\theta^+(x_t|\cdot)$ and $p_\theta^-(x_t|\cdot)$ are attribute-conditional LMs for the desired and undesired attributes, respectively, and ω is a weighting hyper-parameter. This method corresponds to our framework with the following energy function:

$$E_\theta(c|x_{1:t-1}) := -\omega \Big(\log p_\theta^+(c|x_{1:t+s}) + \log p_\theta^-(c|x_{1:t+s}) \Big) ,$$

B Other Related Prior Work

Our framework takes inspiration from prior methods that attempt to directly sample from the joint EBM in Eq. 2. For example, Khalifa et al. (2021) propose to learn another LM $\pi(x)$ to approximate the joint distribution $p_\theta(x, c)$ such that we can directly sample from $\pi(x)$. Other works (Bakhtin et al., 2021; Bhattacharyya et al., 2021) use the so-called self-normalized sampling, which proceeds by first sampling a set of samples from $p_\theta(x)$, and then resample (rerank) using $p_\theta(c|x)$. Kumar et al. (2021) propose to sample text by a running a continuous optimization that resembles the Langevin dynamics (Welling & Teh, 2011). Another line of works focus on fine-tuning LMs for CTG. Shirish Keskar et al. (2019) fine-tune a conditional LM using a variety of control keywords as the prompt. Chan et al. (2021) extend Shirish Keskar et al. (2019) to enable using phrases as the control prompt. Chan et al. (2021) propose a latent space generative model with an encoder-decoder framework. Cao & Wang (2021) fine-tune an LM for style-controlled summarization generation. Ross et al. (2021) train a T5 model (Raffel et al., 2020) on a custom corpus for a given task.

Table A1: Applicability of baselines to various CTG scenarios

Baselines	Compatible Experiment Settings				
	single word	topic	sentiment	Detoxification	multiple control
REINFORCE	✔	✔	✔	✔	
REINFORCE_Px	✔	✔	✔	✔	
PPO (Ziegler et al., 2019)	✔	✔	✔		
GDC	✔	✔	✔	✔	✔
CTRL		✔	✔		
DAPT			✔	✔	
PPLM	✔	✔	✔	✔	✔
GeDi			✔	✔	
FUDGE	✔	✔			
DExperts			✔	✔	
PnB					✔
Ours	✔	✔	✔	✔	✔

These methods require either extensive training/fine-tuning for each single attribute or expensive sampling procedure, making them less useful in practice for flexible, efficient, plug-and-play CTG use cases. For example, GDC (Khalifa et al., 2021) takes more than 100 hours to train on Quadro RTX 8000 GPUs. In contrast, our method requires no training/fine-tuning.

There also exist a number of methods that only focus on a specialized CTG scenario. For example, Ross et al. (2021) study text rewritting, and Ghazvininejad et al. (2017) specialize in poetry generation. Notably, Liang et al. (2021) focus on understanding and mitigating the biases in LMs. Liu et al. (2021) fine-tune pre-trained LMs for emotion-controlled text generation. Different from the above methods that are either only applicable to or evaluated on the limited CTG use cases, our approach is flexible for a wide range of commonly studied CTG problems.

C Experiment Details

In this section, we include additional experiment details.

C.1 Configurations of Our Method

For the single word control and topic control experiments, we set $\lambda = 1$ for experiments in Table 2. For analysis, we vary the hyper-parameter from 0.1 to 10, with steps of 0.1 from 0.1 to 1 and steps of 1 from 1 to 10. For clarity of presentation, we report results of selected values in Figure 1 because the conclusion remains unchanged. For the sentiment control experiment, for positive sentiment, we set $\lambda = 15, \tau_c = 0.995, \tau_{\bar{c}} = 0.001$. For negative sentiment, we set $\lambda = 10, \tau_c = 0.999, \tau_{\bar{c}} = 0.001$. For the detoxification experiment, for both toxic and non-toxic prompts, we use $\tau_{\bar{c}} = 0.1$, i.e., we avoid all tokens above toxic probability of 0.1, as predicted by the perspectiveAPI. For the composition of controls experiment, we use $\lambda = 5$ for topic control and $\lambda = 10, \tau_c = 0.999, \tau_{\bar{c}} = 0.001$ for both positive and negative sentiment controls.

Choosing keywords. Here, we clarify how we choose keywords for each control attribute under consideration in this work. For topic control, we use the same keyword lists associated with each topic as in (Dathathri et al., 2020). For sentiment control, building the keyword list proceeds as follows. First, we select an external classifier, which in our case is the Huggingface sentiment classification pipeline (Wolf et al., 2020). Second, we select a threshold for positive and negative sentiment, respectively. The selected thresholds are discussed in the paragraph above. Third, we loop through each word/token in the base LM's vocabulary. If the word/token falls within the threshold for positive (negative) sentiment, we add it to the keyword list associated with positive (negative) sentiment. For detoxification, the procedure is almost identical to sentiment control, except that here, we use the toxicity classifier from Perspective AI (Perspective, 2021).

C.2 Configurations of Baseline Methods

We largely follow the configurations from the official implementations of the respective baselines. Below, we detail the setup for each baseline in each experiments. Unless otherwise stated, the base language model is a small (12-layer) GPT-2.

GDC, REINFORCE, REINFORCE_Px, and PPO. We use the implementations[1] and configurations of Khalifa et al. (2021), specifically, Table 5 in Section F in their Appendix. Note that some configurations provided in the code repository[2] do not match with the ones suggested in the paper; in those cases, we try our best to change and match the configurations that appear in their paper. GDC is used in all experiments while the other three baselines are used in all but the composition of multiple control experiments.

PPLM. We use the official implementation of Dathathri et al. (2020) and the suggested configurations therein.[3] PPLM is used in all experiments. In particular, for the single word and topic control experiments, we use the following configurations:

```
--length 40 --gamma 1.5 --num_iterations 3 --num_samples 100 --stepsize 0.03
--window_length 5 --kl_scale 0.01 --gm_scale 0.95 --sample --colorama
```

For the sentiment control experiment, we set `class_label=` 2 for positive sentiment and 3 for negative sentiment, along with the following configurations:

```
--length 40 --gamma 1.0 --num_iterations 10 --num_samples 100
--stepsize 0.04 --kl_scale 0.01 --gm_scale 0.95 --sample
```

For the composition of multiple controls experiment, we use the following configurations:

```
--length 60 --gamma 1.5 --num_iterations 5 --num_samples 500
--stepsize 0.03 --kl_scale 0.01 --gm_scale 0.95 --colorama --sample
```

For the sentiment control experiment, PPLM requires a larger GPT-2 (24 layers) as the base language model to be compatible with the sentiment model they used which is a 24-layer GPT-2 in PPLM. For all other experiment, we use the same small GPT-2 (12 layers) as the base language model in PPLM to be consistent with other methods under comparison.

Fudge. We use the implementation and the configurations of Yang & Klein (2021).[4] We change the `condition_lambda` hyper-parameter from the default value 1 to 4 as suggested in the README in their code repository. We following the experiment settings in Yang & Klein (2021) and use Fudge in the single word and topic control experiments in our paper.

GeDi. We use the implementation of Krause et al. (2020) with the default configuration in the code repository.[5] Following the experiment settings in Krause et al. (2020), we use GeDi in the sentiment control experiment in our paper.

Plug-and-blend. We use the implementation of Lin & Riedl (2021) with the default configuration. This method extends GeDi to enable control over multiple topics. We have attempted to adapt this method for multiple topics *and* sentiment controls but it generates the same samples regardless of the sentiment. We have also attempted contacting the authors and will add the topic-and-sentiment control results once we hear back from the authors' suggestions. For now, we use Plug-and-blend in the composition of multiple controls experiment in which the controls are combinations of topics without sentiment.

Dexperts, CTRL, and DAPT. We use the implementation of Liu et al. (2021) for Dexperts and their adapted implementations of CTRL and DAPT which unifies the configuration,[6] making it convenient to run and reproduce results. For CTRL and DAPT, we use the default configurations. For

[1] https://github.com/naver/GDC
[2] We use the configurations in https://github.com/naver/GDC/tree/master/configs
[3] https://github.com/uber-research/PPLM/tree/master/paper_code
[4] https://github.com/yangkevin2/naacl-2021-FUDGE-controlled-generation
[5] https://github.com/salesforce/GeDi
[6] https://github.com/alisawuffles/DExperts

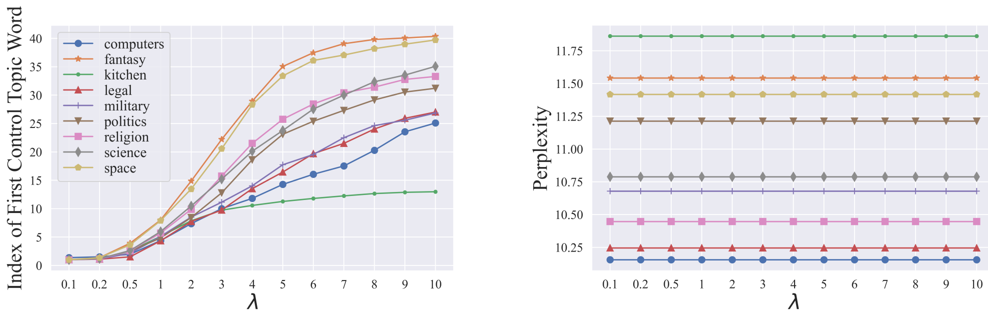

Figure A1: Analysis of how the hyper-parameters λ in Eq. 12 influences control satisfaction (**Left** plot) and language quality (**right** plot). We see that λ controls how early a keyword in the topic keywords list appears in the generated text but does not impair language quality.

Dexperts, we use the example configuration provided in the README in their code repository. Following the experiment settings in the respective papers of these methods, we use all three baselines in the sentiment control experiment and additionally use CTRL in the topic control experiment. Except for CTRL which only provides the largest GPT-2 (48 layers) pre-trained model, we use the small GPT-2 for Dexperts and DAPT. Note also that CTRL requires a special "control code" prepended at the beginning of any prompt as a way to impose control during generation. The control codes required to achieve topic and sentiment control are summarized in Table S7 in Section S7.1 in the Appendix of Dathathri et al. (2020). Because only five out of nine topics we used are compatible with the CTRL setting, we only evaluate CTRL on the subset of these five topics in the topic control experiment.

C.3 Human Evaluation Settings

We perform human evaluation via Amazon Mechanical Turk (MTurk).[7] for the sentiment controlled generation and detoxification experiment. The evaluation proceeds as follows. First, we select highly competitive baselines from all methods based on their performance on automatic evaluation metrics in addition to our method. For sentiment controlled generation, we select GeDi, DAPT, PPLM, CTRL, PPO, GDC and DExperts. For detoxification, we select PPLM, GDC, DExperts, GeDi, DAPT, and PPO. We randomly sample 100 generations from each method for each controlled setting, which include positive sentiment, negative sentiment, and detoxification. Thus, in total, for sentiment controlled generation, we evaluate $100 \times 8 \times 2 = 1600$ generated samples and for detoxification, we evaluate $100 \times 7 = 700$ generated samples. Second, we determine our evaluation criteria. For sentiment controlled generation, we ask human evaluators to rate on sentiment on a scale from 0 – 4 ranging from "very negative", "negative", "neutral", "positive", and "very positive" and on fluency on a scale from 0 – 2 ranging from "not fluent", "somewhat fluent", and "fluent". For detoxification, we ask human evaluators to rate on toxicity of the *continuation* generated by each method without considering the prompt on a scale from 1 – 3 ranging from "non-toxic", "toxic" and "very toxic" and on fluency with the same setting as in the sentiment controlled generation. Third, for each generated sample, we ask three different evaluators to provide their ratings on our evaluation criteria, leading to a total of $(1600 + 700) \times 3 = 6300$ individual ratings. We only invite *highly skilled* evaluators to perform our evaluation. These workers are "master workers" on MTurk,[8], have cumulatively completed at least 1000 evaluation tasks on MTurk, and have at least 95% evaluation acceptance rate. Finally, we collect all results and compute the average scores.

D Additional Experiment Results

In this section, we present additional experiment results.

Detailed results on word and topic controls. In tables A3 and A4, we present the complete quantitative evaluation results for each control word and each control topic.

[7] https://www.mturk.com/

[8] https://aws.amazon.com/blogs/aws/amazon-mechanical-turk-master-workers/

Table A2: Language detoxification experiment results evaluated on continuation generations from non-toxic prompts.

Method	Toxicity (↓)	Language Quality (↓) perplexity	diversity (↑) dist-1	dist-2	dist-3
GPT2	0.163	13.630	0.784	0.924	0.984
REINFORCE	**0.074**	20.006	0.648	0.853	0.894
REINFORCE_Px	0.145	18.600	0.755	0.892	0.901
PPO	0.107	13.769	0.788	0.932	0.932
GDC	0.151	13.737	0.788	0.928	0.927
DExperts	0.104	25.250	0.777	0.881	0.873
DAPT	0.125	71.001	0.873	**0.959**	**0.938**
GeDi	0.114	125.815	**0.877**	0.936	0.910
PPLM	0.162	18.085	0.805	0.935	0.929
ours	0.111	**13.625**	0.788	0.928	0.928

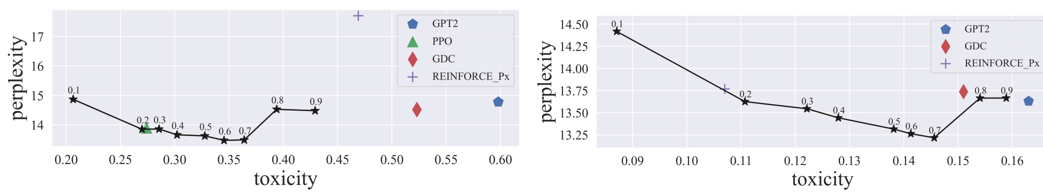

Figure A2: Analysis of how the hyper-parameter λ impacts reducing toxicity and maintaining language quality in the detoxification experiment, for both toxic prompts (**left** plot) and non-toxic prompts (**right** plot). We see that, for various λ's we investigated, our method achieves better trade-off between toxicity and perplexity compared to four strong baselines.

Hyper-parameter analysis for the topic control experiments. We show in Figure A1 how the hyper-parameter λ impacts topic control satisfaction and generated texts' language quality. Here, we say the generated text satisfies the control when any word in keywords list associated with the desired topic appears in the generated text. We observe that, for a more aggressive λ, i.e., when λ is larger, we achieve better control satisfaction because a topic keyword appears earlier in the generated text. For most of the λ values we investigate, the generated texts satisfy the topic control for most of the control topics. We also see that λ almost has not impact on language quality. We hypothesize that because each topic contains a sizable number of keywords, there is a high chance some of these keywords are among the top predictions at each generation step in our base LM. Thus, increasing the chance of sampling these words that are already among the base LM's top predictions will not significantly impact the perplexity scores. Future work will investigate using a combination of keywords, instead of just one, to enforce topic control.

Detoxification results from non-toxic prompts. We show the quantitative evaluation results on generated texts from non-toxic prompts, comparing our method with several baselines. Table A2 presents the results. Overall, these results align with those in Table 4: our method achieves competitive performance in each metric and nicely balances a low toxicity score and high language quality and diversity.

Hyper-parameter analysis for the detoxification experiments. We analyze how the hyper-parameter λ trades off reducing toxicity and maintaining language quality. Figure A2 plots toxicity against perplexity for nine values of λ and compare them to four strong baselines. We can see that, for both toxic and non-toxic prompts, the different λ values achieve better toxicity-perplexity trade-off compared to strong baselines.

Qualitative Results. In Tables A5 – A15, we show a number of generated samples from our method and selected strong baselines for various CTG scenarios.

Table A3: Detailed single word control results for each control word. Our results, not necessarily the best among all baselines, are marked in bold.

Control Word	Method	Control Satisfaction (↑) word accuracy	Language Quality (↓) perplexity	Language Diversity (↑) dist1	dist2	dist3
US	REINFORCE	1.00	666.7526 (2057.1201)	0.314	0.5741	0.7675
	REINFORCE_Px	0.21	24.8845 (31.6511)	0.8328	0.9148	0.8947
	PPO	1.00	9.7831 (3.5822)	0.8256	0.9508	0.9397
	GDC	0.93	9.3510 (2.5348)	0.8186	0.9415	0.9341
	PPLM	0.42	10.4812 (3.6961)	0.8126	0.9368	0.9316
	FUDGE	1.00	11.4657 (3.9645)	0.7725	0.9064	0.909
	ours	**1.00**	**11.3174 (3.3471)**	**0.8246**	**0.9412**	**0.9362**
Canada	REINFORCE	1.00	16.2365 (11.2301)	0.1679	0.2935	0.4031
	REINFORCE_Px	0.33	14.8635 (13.0732)	0.9157	0.9432	0.9124
	PPO	1.00	9.9545 (2.3833)	0.8286	0.9505	0.9407
	GDC	0.84	9.8063 (2.5902)	0.8286	0.9385	0.9334
	PPLM	0.29	10.9400 (3.7051)	0.8138	0.9301	0.9267
	FUDGE	0.26	9.7072 (4.6484)	0.8471	0.9412	0.9277
	ours	**1.00**	**11.6446 (3.3984)**	**0.8221**	**0.9502**	**0.939**
Paris	REINFORCE	1.00	260.0703 (142.3684)	0.3332	0.5992	0.7897
	REINFORCE_Px	0	11.2291 (3.9074)	0.8882	0.9402	0.9106
	PPO	0.99	10.5341 (2.9438)	0.8184	0.9475	0.9359
	GDC	0.91	10.0804 (3.2030)	0.8075	0.9406	0.9351
	PPLM	0.19	10.8496 (4.6177)	0.7974	0.9315	0.928
	FUDGE	0.45	11.9026 (4.7950)	0.799	0.9041	0.8968
	ours	**1.00**	**12.1993 (3.4365)**	**0.8224**	**0.945**	**0.9363**
amusing	REINFORCE	1.00	1616.9402 (1077.0526)	0.487	0.8452	0.8807
	REINFORCE_Px	0.9	26.8899 (11.3656)	0.6284	0.7536	0.7612
	PPO (Ziegler et al., 2019)	0.84	13.7737 (3.8277)	0.8051	0.9488	0.9399
	GDC	0.89	14.1445 (4.0946)	0.8034	0.9426	0.9366
	PPLM	0.02	12.2712 (7.5083)	0.8053	0.9291	0.9268
	FUDGE	0	19.4637 (8.8808)	0.5824	0.7508	0.7907
	ours	**1.00**	**14.1843 (3.4905)**	**0.8246**	**0.9441**	**0.9351**
Wikileaks	REINFORCE	1.00	14.1799 (24.0376)	0.1231	0.1961	0.2665
	REINFORCE_Px	0	15.2540 (5.2523)	0.8081	0.8848	0.8737
	PPO	0.98	10.9353 (2.7979)	0.8278	0.9468	0.9366
	GDC	0.97	10.2984 (2.7075)	0.8376	0.951	0.9401
	PPLM	0.3	12.7863 (5.9673)	0.8416	0.9516	0.9381
	FUDGE	0	11.8344 (3.9445)	0.8231	0.9376	0.9278
	ours	**1.00**	**12.6191 (3.7628)**	**0.833**	**0.9486**	**0.9391**
China	REINFORCE	1.00	21.6350 (14.2200)	0.1933	0.3446	0.4758
	REINFORCE_Px	0.14	30.4373 (21.9684)	0.626	0.7615	0.792
	PPO	1.00	10.0547 (2.3527)	0.8336	0.953	0.9426
	GDC	0.95	9.7273 (2.6467)	0.8317	0.9455	0.9366
	PPLM	0.26	10.4724 (3.0868)	0.8157	0.9355	0.9292
	FUDGE	0.9	11.7443 (3.7337)	0.8098	0.9427	0.9334
	ours	**1.00**	**11.6518 (3.4799)**	**0.8184**	**0.9352**	**0.9273**
amazing	REINFORCE	1.00	250.6540 (189.6618)	0.3128	0.543	0.7278
	REINFORCE_Px	0.5	30.9958 (25.3260)	0.7655	0.8475	0.8419
	PPO	1.00	12.3714 (3.1623)	0.7914	0.9337	0.934
	GDC	0.97	12.4536 (4.2096)	0.8055	0.9318	0.929
	PPLM	0.17	11.4864 (3.3437)	0.8074	0.9373	0.9311
	FUDGE	0.98	9.5816 (5.5922)	0.4321	0.5195	0.5494
	ours	**1.00**	**12.4860 (3.4748)**	**0.8114**	**0.94**	**0.934**
restaurant	REINFORCE	1.00	97.5834 (60.8047)	0.2781	0.521	0.6976
	REINFORCE_Px	0.28	15.6067 (13.5147)	0.8235	0.8737	0.8603
	PPO	1.00	12.3296 (2.9050)	0.8127	0.9478	0.9416
	GDC	0.96	12.3676 (3.1227)	0.808	0.9444	0.9361
	PPLM	0.14	10.9700 (4.0731)	0.7999	0.9223	0.9179
	FUDGE	0.72	13.4720 (3.8018)	0.7879	0.9411	0.9339
	ours	**1.00**	**12.6172 (3.2136)**	**0.8203**	**0.9454**	**0.9396**
Vampire	REINFORCE	1.00	836.7204 (1369.1324)	0.2065	0.3676	0.5007
	REINFORCE_Px	0.85	6.7718 (7.5073)	0.2586	0.2661	0.2627
	PPO	0.99	11.2641 (3.3673)	0.753	0.9031	0.9125
	GDC	0.96	10.4975 (3.6859)	0.7571	0.8848	0.8929
	PPLM	0.12	11.2209 (4.3610)	0.7949	0.9192	0.9172
	FUDGE	0.01	11.4629 (5.0298)	0.7145	0.8381	0.8525
	ours	**1.00**	**13.3584 (3.7057)**	**0.8232**	**0.9437**	**0.9372**

Table A4: Detailed topic control results for each topic. Our results, not necessarily the best among all baselines, are marked in bold.

Control Topic	Method	Control Satisfaction topic accuracy	Language Quality perplexity	Language Diversity dist1	 dist2	 dist3
computers	REINFORCE	1	55.8594 (37.4444)	0.2151	0.3803	0.5281
	REINFORCE_Px	0.51	11.4289 (9.5426)	0.7654	0.8979	0.9003
	PPO	0.99	4.3002 (4.0575)	0.9173	0.954	0.9383
	GDC	0.82	10.0003 (2.9592)	0.7955	0.9231	0.9247
	PPLM	0.64	11.3503 (4.3965)	0.832	0.9476	0.9396
	CTRL	0.51	51.7158 (36.5406)	0.9344	0.9675	0.9427
	FUDGE	0.91	12.5717 (4.2916)	0.7843	0.9242	0.9218
	ours	**1**	**10.1560 (3.6291)**	**0.8052**	**0.932**	**0.9312**
kitchen	REINFORCE	1	3224.4497 (1644.9415)	0.5143	0.8006	0.8718
	REINFORCE_Px	0.26	18.6206 (10.0080)	0.8273	0.902	0.8899
	PPO	0.99	10.8938 (2.7680)	0.803	0.9316	0.9298
	GDC	0.87	10.2482 (3.3945)	0.8077	0.934	0.9288
	PPLM	0.4	10.9309 (3.1395)	0.7976	0.9284	0.9268
	CTRL	-	-	-	-	-
	FUDGE	0.59	16.6660 (5.8355)	0.7939	0.9238	0.9159
	ours	**1**	**11.8615 (2.9079)**	**0.8256**	**0.9477**	**0.938**
military	REINFORCE	1	78.9958 (22.8026)	0.2765	0.6975	0.892
	REINFORCE_Px	0.83	25.8347 (53.6340)	0.9251	0.9298	0.89
	PPO	0.99	9.4962 (2.9360)	0.803	0.9296	0.9232
	GDC	0.84	9.8162 (2.7757)	0.813	0.9366	0.9303
	PPLM	0.61	10.1380 (3.2390)	0.8071	0.9294	0.9266
	CTRL	-	-	-	-	-
	FUDGE	0.93	12.0739 (3.4314)	0.806	0.9405	0.9312
	ours	**1**	**10.6792 (3.1397)**	**0.799**	**0.9361**	**0.9324**
religion	REINFORCE	1	45788.8438 (426901.7812)	0.2352	0.4182	0.5685
	REINFORCE_Px	0.2	14.0497 (9.6284)	0.8198	0.9072	0.8862
	PPO	0.91	8.6393 (2.6998)	0.7853	0.9224	0.9226
	GDC	0.68	8.8640 (3.0208)	0.8193	0.9387	0.9322
	PPLM	0.28	10.1140 (3.4671)	0.8059	0.9315	0.9237
	CTRL	0.4	44.3691 (25.7538)	0.9201	0.9712	0.9488
	FUDGE	0.47	12.2833 (6.3449)	0.7618	0.8703	0.8721
	ours	**1**	**10.4475 (3.1664)**	**0.7959**	**0.9254**	**0.9239**
space	REINFORCE	1	34.5196 (49.2324)	0.1911	0.3258	0.4441
	REINFORCE_Px	0.05	11.7859 (3.4879)	0.8766	0.9368	0.9062
	PPO	0.96	10.2276 (4.2292)	0.7946	0.9368	0.9311
	GDC	0.85	10.7659 (3.4081)	0.7935	0.9301	0.9275
	PPLM	0.25	10.6839 (3.7410)	0.8161	0.9443	0.9346
	CTRL	-	-	-	-	-
	FUDGE	0.65	14.0815 (4.0904)	0.7845	0.9223	0.9254
	ours	**1**	**11.4168 (2.7532)**	**0.8092**	**0.9395**	**0.9308**
fantasy	REINFORCE	1	550.1111 (388.5793)	0.328	0.6009	0.7878
	REINFORCE_Px	0.05	27.8760 (46.9606)	0.8808	0.9388	0.9083
	PPO	0.92	12.5310 (4.4887)	0.7797	0.9263	0.9265
	GDC	0.89	12.1348 (3.8782)	0.7714	0.9146	0.917
	PPLM	0.14	11.0444 (3.5334)	0.8024	0.9269	0.9242
	CTRL	-	-	-	-	-
	FUDGE	0.17	14.7171 (7.0708)	0.8032	0.9264	0.9229
	ours	**1**	**11.5423 (2.9142)**	**0.7945**	**0.9297**	**0.9279**
legal	REINFORCE	1	74.5358 (285.1382)	0.1806	0.3056	0.4105
	REINFORCE_Px	0.37	20.2291 (14.5995)	0.868	0.9235	0.8959
	PPO	1	10.1167 (2.5498)	0.8069	0.936	0.9324
	GDC	0.84	9.3676 (2.8337)	0.8173	0.9447	0.9399
	PPLM	0.61	10.4652 (3.2302)	0.81	0.9402	0.9318
	CTRL	0.39	40.7088 (28.9848)	0.9243	0.9742	0.9499
	FUDGE	0.9	10.2886 (2.7878)	0.8186	0.9466	0.9369
	ours	**1**	**10.2460 (2.8621)**	**0.8085**	**0.9391**	**0.9357**
politics	REINFORCE	1	225.6859 (335.7491)	0.2874	0.6014	0.8273
	REINFORCE_Px	0.62	27.3640 (16.5313)	0.666	0.7916	0.8131
	PPO	0.98	10.6021 (2.6582)	0.8133	0.9433	0.9337
	GDC	0.88	10.0484 (2.8276)	0.8154	0.9443	0.9367
	PPLM	0.34	9.9450 (3.2881)	0.8152	0.9335	0.929
	CTRL	0.34	48.1547 (33.1596)	0.9384	0.9709	0.9464
	FUDGE	0.78	10.7969 (4.0509)	0.8205	0.9439	0.934
	ours	**1**	**11.2136 (2.6106)**	**0.8158**	**0.9449**	**0.94**
science	REINFORCE	1	581.6110 (348.2726)	0.4311	0.7351	0.854
	REINFORCE_Px	0.15	25.9675 (25.1919)	0.7648	0.8534	0.8459
	PPO	0.96	10.6762 (3.5327)	0.806	0.9373	0.9334
	GDC	0.8	9.4292 (2.6120)	0.8108	0.938	0.9323
	PPLM	0.33	10.5173 (3.1171)	0.824	0.9463	0.9381
	CTRL	0.32	58.5154 (60.7074)	0.9307	0.9674	0.9402
	FUDGE	0.77	12.6211 (3.2532)	0.8009	0.9387	0.933
	ours	**1**	**10.7880 (3.4760)**	**0.8051**	**0.9351**	**0.9295**

Table A5: samples of generations in the word control scenario. The control word here is `amazing`.

Method	Generation (control word = **`amazing`**)
REINFORCE	A amazing amazing amazing amazing amazing alliance amazing trees amazing buds amazing trees amazing amazing amazing amazing fruits amazing deltaゼウス amazing amazing trees ゼウス amazing amazing amazing amazing trees amazing amazing pokemon amazing pokemon amazing omegaKo incredible alliance
	The amazing amazing items amazing apples amazing fruits amazing trees amazing fruits amazing trees amazing flowers amazing amazing fruits amazing Oakland amazing Emeraldゼウス amazing trees amazing amazing Emerald amazing amazing amazing amazing delta delta amazing Oakland amazing amazing omega
	The remarkable amazing orbs amazing amazing amazing trees amazing amazing amazing amazing trees amazing amazing fruits amazing amazing fruits amazing trees amazing fruitsELL incredible amazing days amazing amazing delta amazing amazing pokemon amazing amazing amazing omega amazing omega amazing
	Welcome amazing orbs amazing amazing fruits amazing amazing flowers incredible amazing amazing trees amazing amazing flowers incredible amazing delta amazing amazing days amazing Oakland amazing ib amazing amazing Emerald Awakening amazing Oakland amazing omega incredible amazing amazing Emerald amazing days
REINFORCE_Px	A new study shows that women're amazing to have around them. Researchers continue to share amazing news and amazing stories to share them
	Say thanks by giving a tip and helping them continue to share amazing Things with the Thingiverse community. We're sure
	A community group has launched a petition calling on the community to share a moment and community moment together to celebrate the amazing community community
	We're sure them to share this amazing news late this morning with the community thanks to a quick share by our community amazing community
PPO	There's a huge amount going on. And it's not only because of amazing technology, but because of the amazing people who are there to create this content, and to help out. The creators are
	It was amazing seeing how far you could go in your life. It's hard to think of a better feeling than feeling like a person. This article will show you how to make it.
	The most amazing thing I've ever done for my birthday is to do it with a gift! It's so easy! The only thing I've done that I have to mention is my amazing friends,
	The amazing thing about the new iPad has been how amazing it really is. I've been using the device for nearly three years now and have found it to be so easy to use that when you are
GDC	The most amazing thing about this game is the way it makes you feel. If you are feeling hungry, hungry for something, hungry for some kind of experience, you will want to check out this game
	The new version of the Windows Phone 8 software comes with an amazing new feature that you can't get from previous versions of Windows: it's called Windows Defender. Windows Defender is an amazing new
	In this post we'll look at how to create a simple, reusable, and amazing web framework from scratch. In order to accomplish this, I've had to spend some time creating my own custom framework
	It is amazing that we have not been able to see any of these other amazing projects, but I can't think of a moment in my life where a project like This is the New York Giants is
PPLM	We are a small, yet very successful company that has developed a lot of great products. We've built our reputation on being very well rounded and very friendly. We've had a
	The world's first real-life virtual reality platform is coming to PlayStation VR, with the release of the PlayStation VR version of Half Life 3 in 2017. We've been working hard on Half
	"You have to look at what's happening in China and the other world," he tells me. "The people have no clue about the real value of the world." "And it's
	As the year of the new millennium has passed, and the world has changed in some amazing ways, some amazing amazing things have happened. The new millennium is here. We are living in the new millennium
FUDGE	We celebrate the incredible, amazing, beautiful, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing
	Hey guys, I've been working on this project for about a week now, so I figured it wouldn't be amazing. I've been working on the new system for a while now and
	Every great artist is amazing. Every amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing, amazing
	Share this: Pinterest Email After winning the IndyCar series for the first time in his life in 2011, Lance Stroll has made incredible, amazing, incredible, amazing, amazing.
ours	The United States is on track to meet its international commitments on climate change by 2022. The amazing thing is that we are already on the cusp of this goal. The United States is now the world
	As we approach the year 2023, one of the amazing things that happens in life is the change that happens when you get older and start seeing the world differently. It is the age of the Internet
	We're excited to be announcing our first ever amazing partnership with a local artist, the artist's daughter, and the artist's son. We've had a lot of fun collaborating with our local community,
	The following article is a transcript of a conversation with amazing people at a local restaurant in the city of Los Angeles, California. "Hi there, this is my friend, I just wanted to

Table A6: samples of generations in the word control scenario. The control word here is `restaurant`.

Method	Generation (control word = `restaurant`)
REINFORCE	restaurant restaurant, a restaurant restaurant Choi restaurant restaurant China Myanmar Myanmar Jian restaurant restaurant The restaurant restaurant ChengTangTang restaurant restaurant restaurant Jian restaurant restaurant India restaurant restaurant restaurant restaurant China restaurant restaurant restaurant restaurant restaurant restaurant China
	SAN Restaur restaurant the restaurant restaurant restaurant restaurant Choi Myanmar restaurant Jian restaurant restaurant Choi restaurant restaurant restaurant Jian restaurant restaurant restaurant Myanmar Choi restaurant Indian restaurant India restaurant restaurant restaurant restaurant restaurant restaurant restaurant restaurant restaurant restaurant Cheng restaurant
	Abstract Cook restaurant restaurant Cheng restaurant restaurant restaurant restaurant KKI restaurant Panda restaurant Myanmar restaurant restaurant restaurant Panda restaurant restaurant restaurant restaurant restaurant restaurant restaurant Bangladesh restaurant restaurant Choi restaurant Panda restaurant restaurant Indian Indian Restaur restaurant
	The restaurant restaurant Jian Liang restaurant The restaurant restaurant Choi restaurant restaurant Myanmar Cheng China restaurant The Seoul restaurant Choi restaurant restaurant restaurant India Indian restaurant restaurant China restaurant restaurant restaurant restaurant restaurant Indian Restaur restaurant Bangladesh India restaurant India
REINFORCE_Px	A man died for his service to an American hero last week. James Holmes killed four people at a gun range this
	In the midst of its worst economic downturn in decades, China faces a new threat to its security that has shaken its traditional allies
	In this Nov. 3 restaurant illustration released Tuesday, Jan. 1 Yeshiva University students offer $10 checks during a Yes
	Does this restaurant offer table service? Yes No Unsure Is this restaurant good for local cuisine? Yes No Unsure
PPO	I'm working at a restaurant that specializes in the kitchen. It's a big deal for me. I get called to the door to see if there's a customer in the kitchen and they'll say
	A restaurant in the city's east end has been closed after two people were shot and killed while they attempted to help the injured. The shooting began at 7:45 p.m. Sunday
	A New Yorker restaurant has been hit by a car driven into a building, sending it flying into a building in the heart of San Francisco. The car crashed into the restaurant on the corner of
	The world's biggest restaurant chain is now making a big deal about its "fast, easy, fast" policy. The company, based in San Francisco, plans to make $10,000
GDC	The city of Los Angeles is planning an expansion of its popular restaurant, the Tenderloin, on Westheimer Avenue, which is expected to open in 2017. The project is the largest
	The restaurant, which will feature a full-service menu in the coming weeks, is a new, boutique boutique restaurant in the trendy downtown Toronto neighbourhood of Oakville, Ont. The concept, which will
	The restaurant is currently closed to the public. We're sorry, currently this live video stream is only available inside of Utah or an approved RSL broadcast territory. We base your location on your
	In this episode, we explore the restaurant scene at the end of the day. The restaurant business is a booming market for many reasons. One is that it has been in business for centuries,
PPLM	In the first round of the 2014 NFL Draft, the Houston Texans went through two of their best prospects in receiver Antonio Brown. In the second round of the 2014 NFL Draft, the Houston Texans
	It's been an exciting year in the NBA's new era. The Nets' franchise is back in action, but what does that mean for the team? The Nets will have to make some
	I am not sure if you know how much I love my new favorite restaurant, and I am glad we are not the only restaurants serving the restaurant that we are serving. I have been dining here for
	A group of people has been arrested in a shooting attack in the city's Chinatown neighborhood where a group of Asian immigrants is being questioned by police, police chief said. The shooting happened just outside
FUDGE	About: My mom bought this from my mom in a restaurant in a small town, California. She's been a restaurant owner for a while now. I bought this when I had an accident
	About 12 minutes into dinner, a restaurant owner called me up and asked me to come back to the restaurant. I did. The next day, we got into a heated discussion on the restaurant
	About 1,000 people walked away from Starbucks, McDonald's and other fast-food restaurants from a protest in Chicago that drew about 1,000 people who had gathered on the street outside the restaurant.
	About a month after the Boston Marathon bombing that left three people dead and at least 25 injured, Boston Police Superintendent Ray Kelly is on a hunger strike after he said he would bar anyone at restaurants, bars
ours	The following is a guest post by Mark Roussel, restaurant manager at Biscuits and Pastries. Mark has been a member of Biscuits, and was the founder of The Bisc
	A couple of weeks ago we wrote about how restaurant owners are trying to figure out the best way to keep customers happy and healthy. Now, a new survey from the American Council on Exercise finds that a
	In the late 1990s and early 2000s, restaurant chain McDonald's opened a small restaurant in New York City's Chelsea neighborhood with the intention of expanding into the larger East Village. As the chain grew
	AUSTIN, Texas – A restaurant owner in Dallas has been fired after he was caught using his cellphone in a restaurant parking lot without authorization. "I have a lot of respect for the

Table A7: samples of generations in the word control scenario. The control word here is `amusing`.

Method	Generation (control word = **amusing**)
REINFORCE	The amusing pairatorium amusing barrel amusing tree amusing frogTea amusing barrel amusing woodlandTea amusing woodland amusing bed amusing tree amusing gardenTea amusingTeaJew amusing stenLegend amusing woodland amusing sten amusingTea amusing woodlandTea
	Dear amuse amuse amusing bird amusing barrel amusing tree Tree amusing garden.Tea amusing woodland amusing sten amusing creek amusing sten amusing creekMrLegend amusing woodland amusing creekTeaTeaTea amusing amusing treeTea amusing tree amusing
	amusing pair amusing pair amusing barrel amusing garden amusing tree amusing tree amusing treetletletle amusing trees amusing tree brilliants amusing creekometimes amusing trees redes amusing tree amusing sinking amusing woodland amusing amusing stormtroopers
	Dear amusing tree amusing barrel amusing barrel.JerryWritername amusing frog amusing woodland amusing birdBirdTea amusing frogtle amusing woodlandTea amusing sten amusingTeaTea amusing barrel amusing trees amusing amusing woodland amusing amusingTea
REINFORCE_Px	If there was one thing that was universally amusing about the NFL was that it was universally amusing about the NFL was that it was
	If there was one thing that was universally amusing about the NBA was that there was one thing that was universally amusing about the NBA
	If there was one thing that was universally amusing about Marvel's Marvel Cinematic Universe was that it was universally amusing about Marvel Cinem
	The NFL is one of the most popular teams in professional sports during its first decade on television. If there was one
PPO	The first time I heard of this idea I was surprised. I thought it was amusing and I liked the idea, but when I read more about its potential it seemed like a waste of time.
	"We've always had an amusing story about the life of a child," says the writer, who was born with the condition. "There are things about the world that seem more than amusing in their
	The world's first fully functional mobile gaming console, the Sony PlayStation 4, is being developed for the PC by Sony Computer Entertainment America. The Sony Playstation 4 will run a version of PS4
	An amusing thing that has happened in the last couple years has been the emergence of new forms of communication. In the last decade, there have been several new ways of communicating and exchanging information, and there
GDC	The amusing thing about a lot of this kind of stuff is that it usually involves a lot of self-congratulatory and amusing comments. Some of them are downright amusing in nature, like this one
	The most amusing, most shocking thing in the world isn't how much money the government spent on a war, but how much money the public spends on war. Here's how the US government
	"You've got to have something amusing about you," Trump said on Twitter. "I love my country. I hate the people that live in my home country." He was referring to Trump
	It was amusing to think back to the time when a couple of our friends in the area were in need of an amusing way of getting to know each other and the people we were with. The first
PPLM	A man who claimed the death of three women at a London hotel was a 'bitch' has revealed he's a'sadistic idiot'. He also admits to having a'silly'
	In the past year, the company has been working with the government on the idea of making it easier for the public, in part by adding the following to its blog: "We've had
	The new book "What Happened To The White House in the 1970s?" is one of the most interesting books I've read recently about politics and politics of the late 1960s and early 1970s
	I have no idea when the "F" word was invented, but this post contains links to some of my personal experiences with the word in print. I remember a story that was printed in
FUDGE	It was a bizarre and startling experience. I've been a guest of an unexpected guest, and I've found a way to make a very bizarre, very bizarre, very bizarre, very bizarre, very
	He is fond, fond of a lot of things. He has a funny accent, and has been known to be a bit of a prick, so he is very fond of his own penis. He
	I've been noticing and criticizing, especially since my post on this one in the recent past, a subtle, subtle, subtle, very subtle, subtly subtle, very subtle, very subtle, subtly subtle
	I've always wished, whenever a story has an unspoken hint, a hint, a hint, a hint, or a hint, the reader would be more interested in seeing the reader's imagination blown
ours	The following is a list of the most amusing things you can do to make your dog happier. I know it sounds a bit like an easy task, but it really isn't and I am
	"We've got a problem here that needs to be solved," amusingly stated a reporter on a recent day. "And I don't want to get into any of the details, it's too
	This post is part 2 of a series in which I will show amusing stories about how our world is going backwards. This time around, I'm going to show that you can make a difference in your
	The new year has been a very busy one for me! As amusing as it has been to say the least, the holidays are here and it seems like a great time to celebrate! My Santa was

Table A8: samples of generations in the word control scenario. The control word here is `Vampire`.

Method	**Generation (control word = `Vampire`)**
REINFORCE	The Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire III Vampire Vampire Vampire Vampire Vampire Vampire Vampire I Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire The Vampire Vampire Vampire Vampire
	V Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire The Vampire Vampire Vampire, and Vampire Vampire Vampire Day Vampire Butcher Vampire Vampire
	A Vampire. Vampire Vampire Vampire Vampire Vampire Vampire Vampire Vampire. Vampire Vampire Vampire Vampire Vampire
	The Vampire Vampire Vampire Vampire Vampire Red Vampire Vampire Vampire Vampire Vampire Vampire Vampire III Vampire Vampire Vampire Vampire Vampire Vampire The, Vampire, Vampire Vampire Vampire, Vampire Vampire, Vampire Vampire Vampire Vampire II
REINFORCE_Px	Creature — Vampire Wizard
	Creature — Vampire Wizard
	Creature — Vampire Wizard
	Creature — Vampire Wizard
PPO	The Vampire Slayer's most recent film, Vampire Hunter, opens to rave reviews after its release. The film, in which the titular vampire is hunted by his fellow vampires by a mysterious force
	Vampire's Tale 2 by A Vampire's Tale is a series in which two young vampires (Bobby and Susan) are both killed by an unknown assailant after a while. In
	In this episode, Buffy the Vampire Slayer writer and co-writer Matt Bomer shares all the details about what's going on in the world of Buffy the Vampire Slayer. He also has a few thoughts
	The Vampire Slayer was the first major studio album to be released commercially in the United States, making it the first major hit of 2000, and a major influence on the rise of the pop genre. The
GDC	Buffy the Vampire Slayer Season 1 – Season 2 Buffy the Vampire Slayer Season 1 – Season 2 Buffy the Vampire Slayer Season 1 – Season 2 Buffy the Vampire
	Buffy the Vampire Slayer: The Movie Buffy the Vampire Slayer: The Movie The film is an adaptation of the popular novel by M. Night Shyamalan and directed by David
	Vampire's Blood is an episodic adventure RPG with an emphasis on exploring ancient lore and exploring the supernatural and ancient worlds of the Elder Realms. The main characters will explore an alternate world of the Vampire
	The Vampire Lord has been in the headlines recently for a series of disturbing and bizarre incidents at various clubs in the city of Los Angeles. The first was reported by The Vampire Lord, a website
PPLM	The New York Times's John Pilger, a Pulitzer Prize-winning journalist, recently wrote a book titled "The New York Times: A History of Its Own Making." In it, he tells readers
	The United Methodist Church has announced that it's going to donate $2 million to the "Dixie Project" and its local chapter, which has been the inspiration behind a number of controversial social justice initiatives
	"It is a shame that a number of the young men who participated in the protest, who came forward to tell their stories, were not included as well as those who took part in the 'Black
	"You have to look at what's happening in America and the other world around you," he said. "There have been a number of countries where the economy is growing faster than the people are doing
FUDGE	Image Comics #0 (of 4) "A Hero's Tale" Writer: Joe Johnston Artist: Dan Slott Cover Artist: Tom Brevoort
	JOHANNESBURG, South Africa (Thomson Reuters Foundation) - A German-born UEE researcher named "Mortgagehog" faces extradition to South Africa, according to
	Chapter Text Samus's heart was pounding with anticipation. Samus's arms folded, the arms spread wide, and the breaths the rhythmic beats from her heart. A
	Kanye "Kanye, Kanye, Kanye, Kanye" by Chris Graydon Kanye "Kanye, Kanye" by Chris Graydon The Last Knight
ours	I am very pleased to present the final draft of the Vampire Diaries comic, "The Vampire Diaries: The Complete Story," by David Kincaid. The comic is the third of two,
	The new "The Last Vampire in Town" has been released on Blu-ray and DVD. The movie follows a group of vampires who were murdered by a gang of vampire hunters while attempting to save their
	I'm in a room in front of an open doorway, and I see Vampire Lord Ryloth, sitting in the chair with her hands raised, sitting on a small chair. Vampires
	We're excited to be announcing our first ever "Merry Vampire" event, where you can meet the talented Vampire and find out what's coming up in the next season finale.If you

Table A9: samples of generations in the topic control scenario. The control topic here is science.

Method	Generation (topic = science)
REINFORCE	Data data McGill data GowData data data Hercules data Genesis data data BlackRuby data data JonasThor data RubyPython data GenesisSQL data JonasEvaPokemon data data data BlackRubySQL RubyData dataAliceData
	University McGill data data BlackData data data data GenesisRomanDataRuby data data data AliceJones data data BlueRubyJava dataPokemon data BlackDATAData data data BlackArkSQL data AlicePokemon data dataRuby
	Data data data Data KalazelPokemon data Hercules data Hercules data data Hercules data data data Ruby data data Corinth CorinthPokemon data data Kalorean data data Bluegray data CorinthPokemon data AngelenosData data Corinth
	Abstract data Raw data Krishna Krishna data Genesis data data Gabriel data data Alice data EdenBlueSQLSQL data data AliceAlice data data Black dataPokemon dataPokemon data data data data dataArk data data Alice data
REINFORCE_Px	We have a few discounts on the latest discounts on discounts on pricing pricing discounts on pricing discounts on pricing discounts on pricing
	A man accused of sexually assaulting a young girl on a bus in the Bronx is suing the woman who assaulted him.
	It is illegal in the US to sell the use of medical marijuana on the black market. You will see discounts discounts discounts.
	The best of Bill Belichick, Brian Dawkins, John Schneider, Mike Tomlin, Bill Belichick, Bill Belichick, Bill Belichick,
PPO	The National Science Foundation and the Smithsonian Institution will be working to increase the scope of research and the scope of research on the subject of climate change in future generations, according to a letter sent to NASA.
	The US has made a big step forward in its effort to combat online piracy, according to research conducted by the Electronic Frontier Foundation. The US Department of Commerce released its latest online protection policy in
	”The data we have now indicates that there may be a lot of potential for the United States to become a global energy superpower,” said John W. Campbell of the University of Illinois at Urbana
	The US and Canada have reached a historic agreement to set up an International Space Station for human research and exploration. This would be a step towards a new era of space tourism. This is a
GDC	A new study shows that a variety of diseases can result from a combination of genes and environmental factors, with the most severe cases being associated with genetic disorders in the offspring of those with the most severe cases
	”This is not a matter that I want to discuss at all,” said the man, who declined to be identified. ”It’s just something that happens to me all the time, and that I
	It’s no secret that the internet is a huge part of the American culture. A survey last year by The Washington Post found that about half of respondents said they’d seen an Internet connection during
	We’re going to be talking about the fact that we are not a nation of scientists, but of the people who are scientists. We are not just an experiment in a laboratory. We are the people
PPLM	In order to make sure that your system can run on the same system as your PC, you must configure the kernel for your Linux system. The system will be used in the configuration file as part of
	The United States is now the world’s largest consumer of renewable energy, but the U.S. has only one of the most efficient fossil fuel-powered power plants. It’s a big
	In his recent book ”The End of the End” Michael Pollan writes that there is still ”no evidence” that the US government is using the nuclear codes to develop a nuclear weapon.
	The latest news, analysis and opinion from The Daily Telegraph, The Financial Times and The Daily Star is based on research from the International Monetary Fund, the International Monetary Fund, the Institute for Economics and Statistics
FUDGE	Dawn, in a recent paper, has found a new type of DNA to help explain the origins of the modern world. A new gene found in mice has the potential to explain why humans evolved to
	The following article contains an extensive review of the theory and practice of electro-magnetic and magnetic (EMG) magnetic field (MRM) field. This review discusses the basic theory of MRM
	Juan Mata, the world’s greatest super-spender, has been forced to apologise after a row over an unsolicited £100,000 offer, saying the amount he would have made for an
	The world’s most powerful scientists will soon discover that the universe contains a single molecule—one that’s a single atom—that may contain the key to making life possible. The discovery could lead
ours	”It’s not as though climate change has been a major issue for the US,” said the US Geological Survey’s director of the climate and land sciences, Robert Greenfield. ”The US hasn’t
	The latest installment of the data-driven data mining and data mining community, Data Science News, has come to a close. As the data-driven community gathers data, it also publishes its findings.
	We all know that energy consumption is the main determinant of health, as is the amount of energy in our bodies, including blood, skin, bones and teeth. A study published in 2014 in the
	A new study finds research suggesting an increase in autism spectrum disorders is not a new phenomenon. Researchers from the Johns Hopkins School of Public Health and the University of Pennsylvania found the autism prevalence in children

Table A10: samples of generations in the topic control scenario. The control topic here is space.

Method	Generation (topic = space)
REINFORCE	SCPSCP space space space space space Space space space space Mars space space space space space space space Mars space space space space space Mars space space Space Mars space space space space space space Canaveral space space spaces Welcome to space space space space space space occupied space space space space space habitatspaceSpace Space space space space space space space space space space space space space space space space space space occupied space space space space SCPSpace space spaces are space space space space space space space space occupied colon space space space space space space space space space space space spaces space spaces space space space space space space Normandy space space space space The space space habitat Normandy Mars space space occupied space space space space spaces space space space Space space space space space space space space space space spaces space Space space space space space space space space space space space
REINFORCE_Px	It's been a busy week. The first two weeks have seen the release of the second instalment of the Star Wars A former police chief has accused the police of trying to "steal" his position as head of the National Commission for Criminal A new poll finds that Americans are divided over a range of issues, some of which are more difficult than others. The world's biggest airline has been fined by US regulators for its "unintended consequences" in the handling of an emergency
PPO	As the world's first ever space shuttle is on the way down to the moon to collect the first samples from the moon, NASA is taking advantage of new technology to help scientists find the source of the In this post, I've been trying hard to find space to share my work on the subject of virtual reality (VR). In my own personal experience, VR is more of a "thing" in A new study suggests that space travelers might find a good place for a trip if they're lucky enough to get to the moon. A new study finds that space travelers might find a good place Astronomy Picture of the Day Discover the cosmos! Each day a different image or photograph of our fascinating universe is featured, along with a brief explanation written by a professional astronomer. 2008 January 25
GDC	A new study finds that children with ADHD suffer from higher risk of developing an autism spectrum disorder. The research, conducted by the University of Wisconsin, Madison, found that children with ADHD, who "This is an amazing moment for me to be able to do this." I've always been a big fan of the "I'm your star" and "The Big Bang Theory" universe The space shuttle Challenger exploded on the moon on Thursday morning. The satellite's crew members were on a mission to collect samples of life-sustaining microbes from the wreckage of the doomed spacecraft The U.S. space agency has confirmed that an asteroid impact has hit the Earth. The asteroid, called Enceladus, hit its orbit on July 9, 2017. It was
PPLM	By David W. Bierstein WASHINGTON, DC – A new federal study found that the average annual temperature of the planet rose in the past two centuries as the oceans warmed and the oceans warmed In recent weeks, the US military has begun to deploy its first unmanned aerial vehicle (UAV) to the South China Sea. The UAVs are part of the US-led coalition effort to It was a big game, so I decided to give it a try. The idea behind it is very simple: a large scale game of "Starcraft 2" with an asteroid in orbit. The The world's first artificial intelligence platform, known as AIB, announced on Sunday it would launch its first satellite in space on a mission in 2020, with the space agency's space agency space agency also
FUDGE	Transdimensional Space, a nonlinear-based simulation of a space environment, has been published in the Proceedings of the National Academy of Sciences this week. The study, which was led by researchers at Stanford By Krista Kostin When NASA took its first look at space in 2013, it was in deep space, but its first mission in the outer reaches of the solar system was the Mars Arguably the most important discovery from the first half of the 19th century was the discovery that the planets in the orbit of Mercury had formed a circle around their stars. The planets formed The Sun was once the largest solar panel supplier in Japan, and it is now the second-largest supplier to Japanese solar systems, after the Japanese solar power plant. The Sun's solar panel
ours	The New York Times recently published a story on space travel, and its implications for humanity. In this post, I'll discuss some of its most significant and controversial points, from its origins to the implications The United States is considering a satellite system to help monitor a possible asteroid strike. The U.S. space agency has already tested the technology, and the U.S. has deployed a A couple weeks ago I started to think spaceflight might be a good time to get a bit closer to the big picture of how things really are for humans and the planets around them. I'm not As we approach the year 2023, starship design is still a hot topic in the galaxy, as well as the future of space travel. In the latest issue of Wired magazine, Chris Roberts talks a

Table A11: samples of generations in the topic control scenario. The control topic here is `fantasy`.

Method	Generation (topic = `fantasy`)
REINFORCE	”Blood vampire Frankenstein vampire”, vampire vampireoldemort vampire prowiatus vampire vampire vampire redes vampire vampire Dracula Frankenstein tissigil vampire tiss vampires vampire vampire Frankenstein vampire vampire vampire withd vampire vampire vampire vampire vampires nomine vampire vampire
	A vampire vampire confronts vampire vampires vampire vampire vampire confronts vampires vampire nomine Frankenstein Frankenstein vampire vampireoldemort Frankenstein vampire vampire withdBuyableInstoreAndOnline vampire vampire vampire Frankenstein vampire vampire Frankenstein vampire vampire Frankenstein nomine vampire Frankensteinoldemortoldemort nomine vampire
	Last night vampire vampire vampire emerges from vampire Dracula arrives vampire vampire vampire Frankenstein vampire vampire vampire vampire weighsBuyableInstoreAndOnline vampire vampire vampire vampires vampire nomine nomine vampire vampire Frankenstein vampire Frankenstein vampireBuyableInstoreAndOnlineBuyableInstoreAndOnlineBuyableInstoreAndOnline vampire vampireBuyableInstoreAndOnline vampire
	Anonymous demon vampire vampire vampire vampire Frankenstein arrives London vampire Cainhyde vampire vampires Frankensteinoldemort vampire Frankenstein tiss nomine vampire Frankenstein tiss vampire redes vampire vampire vampires vampires adapting nomine vampire Frankenstein vampire nomine vampire vampire vampire vampire Frankenstein
REINFORCE_Px	”We’re trying to bring more people to this table to vote.” — Mr. Trump’s initial statement ”This
	The U.S. Supreme Court is expected on Monday to hear arguments over how the Justice Department should proceed with a lawsuit challenging
	We’ll start with the first, and most important piece of information that comes to mind: the information that makes up the first
	The world of Star Trek, the Star Trek movie of the 80’s, may be over, but it will never be the
PPO	As the world continues to grow larger and bigger, a giant black hole is being created at the center of it, a giant black hole being the largest and most important object on the face of the planet
	A giant squid is seen in the sea off Cape Canaveral as a massive, squid-like monster is spotted in the Atlantic. The monster is a large, squid-like, black-winged sea
	The first time a man has been accused of raping a woman he claims was his own, it could prove a monster. The man, who was identified only because the victim was identified on social
	The world’s most famous horror movie is set in this dystopian dystopian future, and the protagonist of the film, a female alien named K.I.G. is tasked by the mysterious alien spirit of
GDC	The New York State Department of Agriculture and Environmental Protection issued a warning Tuesday warning of a possible flood of food-processing plants in North Carolina, citing a lack of proper training for workers who have been training
	A new survey suggests that Americans are less interested in the idea of marriage equality and more likely to believe it was an accident or a giant leap. The poll, which was taken by the Gallup
	The horror of a new study that shows that giant robots may soon be a threat to society’s survival has been revealed. The study published Monday in the journal Science shows, while robots will likely
	The spirit of our time has finally caught up with us. The last time we spoke, on October 23rd, 2013 at the annual New York Comic Con, we had the great pleasure of
PPLM	In the past year, the company has been trying to find a way into its business model without being seen by investors, but with this year it is taking a different approach: it’s trying to create
	A former U.S. Marine who worked as a contractor with Russia on the U.S.’S. military’s ”Iron Curtain” was convicted of treason for allegedly helping to organize anti-
	”We’re in the midst of the most exciting year in our sport,” said Chris Ritter, president and CEO of the National Football League. ”We’ve had a tremendous growth year, but it
	The UESPWiki – Your source for The Elder Scrolls since 1995 This page has been updated since 1995 and contains information on some of the more important things that can be done with this book
FUDGE	Witch Doctor: A Dungeon-Empire RPG is a fantasy roguelike game that uses magic to fight monsters. The goal of the dungeon, however, is to defeat monsters and get rid of
	SAN ANTONIO — One of the man-made disasters of the past couple of years has left a man paralyzed and nearly crippled. He was one of three people whose limbs were severed by
	SAN JOSE, Ca. — When a white male called the shots on one of the most famous white men in the world in the late 1990s, he attacked and robbed a Black woman he believed
	The man wielding a sword is seen in action as a man walks into a bar in London on November 2, 2015. REUTERS/Toby Melville By Simon Jenkins NEW YORK (
ours	”We’ve got a ghost town that needs to be shut down. There’s a lot going on.” ”We’re going to get the hell out of there.” ”We’ve
	I’m so pleased with my new monster. I’m not sure if it was an accident that I’m now playing as a monster or just something that came along after a lot of practice. My monster
	The latest report in the zombie apocalypse is that the American public is going to be living in fear of the coming zombie apocalypse. But what if there were a way to keep people safe? That
	The first time you see an undead creature, you can see it in their eyes, and it’s hard to tell what it is, or what it can see. It seems to be trying to make

Table A12: samples of generations in the topic control scenario. The control topic here is legal.

Method	Generation (topic = **legal**)
REINFORCE	The court court court court court court Court court courtCourt court courtCourt courts court court court court courthouse court court court court court court courtCourt court courtCourt court court Gutenberg Gutenberg court courtIrish courtCourt
	court lawsuit court court court court court court court court court court court court court court church court courts court court court court court court court court court court courts courts court court courts court court court court court
	InCourt court court court court courts court court court church court court court court churchIrishCourt court court court church court Gutenberg court court court court court court court court court court court court court court
	I court court court court court court court court Court court court courtCourt court courts court court court court court courtimer court court courts court court court courts court church court court court courtIrish court court court
REINFORCE_Px	We've shown you some great results using our new 3D Printable Models shown in the video below. The 3
	A few weeks ago, a reader emailed me and asked me to write a post on the recent developments in the ongoing battle of
	The following is a guest post from David Kohn of the Institute for Economic Affairs in Washington DC. The world is
	I love this recipe for a sweet and savory vegan peanut butter treat. If you make it for yourself it will be delicious
PPO	A new federal lawsuit accuses the state of trying to block the passage of a federal law that would have required transgender students to use the restroom that matches their gender identity at school. The U.
	The United Nations will soon be launching a "safe and secure" international law that will be implemented to combat climate change, a top official told Reuters. The UN climate council will set out the
	The Supreme Court has agreed that a California man could be held responsible for his death by having DNA on the crime that led to his arrest. In a unanimous decision last week, U.S
	As the US and China have signed a free trade deal, the US is trying to get its trade relationship with China back on track. But will China be ready if the two countries meet again this fall
GDC	The US Department of Homeland Security on Wednesday announced it had launched a probe into the Trump administration's decision to end the temporary ban on Muslim immigration, the first in decades. The ban was originally
	In a recent opinion piece, the Huffington Post published its own review of the state of the game. It was a thoughtful critique of the state of the game, a critique that would appeal to gamers who
	The first step toward building the most powerful computer ever built, the Raspberry Pi, will start arriving on June 24th at the Raspberry Pi Summit in San Diego, CA. Raspberry Pi is
	A new survey suggests more people in Australia are being told not to buy the country's most expensive food in supermarkets than ever before. The poll found that the average grocery cost in Australia has dropped
PPLM	It was a big game, so I decided to give it a try. The idea behind it is very simple: use a small game that you want to play in a small town in the US.
	The United States government is working on a new law to regulate the sale of marijuana, a major step that could be legal in the states that are already using medical marijuana. "If we legalize
	This week we're joined by the author of the book about the first three decades of capitalism, Charles Murray. He has been in a position of power for the past five decades, and has been a
	The U.S. Navy's Joint Strike Fighter (JS-47) has been deployed to Japan to support the ongoing U.S.-Japan military operation to defend the country's territorial integrity after Japan
FUDGE	"The government should not permit the sale of any unlicensed firearm in the Commonwealth," the bill said. "This is in violation of the law of the Commonwealth of Virginia and violates the Second Amendment right
	The FBI has arrested a former federal prosecutor and the son of a former federal prosecutor in a federal court in Washington, D.C., and he says there is no longer an indictment against him.
	Following the court's decision on Friday, the prosecution will have to prove that the defendants acted in accordance with the law. The court's reasoning would not be binding on the trial court. However, the
	Following a spate of cases of child sexual abuse in India by the alleged perpetrators, India has launched an "indigenous investigation of the perpetrators of child molestation in schools and universities" to ascertain the identity
ours	A judge ruled Thursday that the government can proceed with a lawsuit against an online service provider that was ordered to pay $3.5 million in damages for an attack that left at least seven people dead.
	If you're in a state where marijuana is legal, it can make a big difference on your life. In fact, it's estimated the state's legalization rates have increased by over 40 percent since 1996
	I am a lawyer who is an active member of the National Lawyers Guild. We represent lawyers who are represented by law firms in the U.S. and Europe. My goal is to help
	"The case for the law is strong. It does not require a special interest," he said. "The law should not require you to take the risk of doing anything you don't want

Table A13: samples of generations in the topic control scenario. The control topic here is `computers`.

Method	Generation (topic = `computers`)
REINFORCE	P database database database data database Karn Sind database Bhar database database database database database database database database database database database database database database Kali database database database database database database database database database database database database Karn Karn database A database database database database Black goddess data database database database data database database database database Kali Kali goddess database Oklahoma caste Karn Karn data databases database database Abyss data database database database database Abyss Abyss database database database database database Database database Abyss database data Oklahoma database database database Abyss database database Abyssinia Abyss database database database database database database database database database database database database database Iraq database database database database database Iraq database database database database The database homepage database database the database database database Abyss database data database database database Zion-ipedia database database database database database database database database AbyssDatabase database database database database Abyss database database Abyssdatabase database database database database
REINFORCE_Px	The best news from the first day of school is here. If you're having trouble getting ready to stand out from the other characters in the school's logo or if you wanted to know the A Florida lawmaker will seek to require employers to pay sick leave to sick workers seeking to be fired if they don't have the money to pay the other way. In the state constitution, employees There are currently no articles related to characters with the "Item No. 15901" characters item. If this item Critical Hit Chance from a weapon has +100% critical chance for all enemies The following item requires the following equipment to use correctly. Note: A new item will be created with the item color information updated only once a month in the item description. This item
PPO	Get daily updates directly to your inbox Subscribe Thank you for subscribing See our privacy notice Could not subscribe, try again later Invalid Email This week's top stories are from 11.10.13 - ES News Email Enter your email address Please enter an email address Email address is invalid Fill out this field Email address is invalid Email cannot be used. Try another or register with your social account The Get Newcastle United FC updates directly to your inbox Subscribe Thank you for subscribing See our privacy notice Could not subscribe, try again later Invalid Email Wayne Rooney says he is in love with Manchester United ES News Email Enter your email address Please enter an email address Email address is invalid Fill out this field Email address is invalid Email cannot be used. Try another or register with your social account A
GDC	The U.S. Department of Homeland Security is conducting a national security analysis of cyberattacks in which it's looking into possible attacks by foreign countries, and whether they have the ability to be carried out If I can help you with your application, I will be happy to help you. Please don't hesitate to fill out the form at the bottom of this page and I will try to be as thorough The US is preparing a major overhaul of the country's military, with the military's chief saying the US should take military action against ISIS as soon as it emerges from its initial training camps, according to "I'm so sorry, but this is not a good place to be," he told reporters, adding that he would "not tolerate anyone who tries to use my name." The incident,
PPLM	A few years back I was on this site and was asked to give my opinions on all things Google. I thought that was a good idea, so I started the site. I have since made it In this article In the past, the term "coder" means a programmer who has code written in C. In Python, C++, Python 2.5, and the standard libraries The UESPWiki – Your source for The Elder Scrolls since 1995 This page has been updated since 1995 and contains information on some of the more important things that have happened. The The U.S. Supreme Court on Wednesday ruled that gay and lesbian people can legally marry under a law that is expected to be adopted later this month by states, according to a blog post from U
FUDGE	A new set of maps that gives you an idea of the type of maps and how to use them. A quick introduction on how to use the latest maps in an existing application to create a Here's a quick look at the various options we have in this category:¡—endoftext—¿The 'Gone With the Wind' Project: a 'post-apocalyptic' post-apocalyptic gam By A new class of malware is coming to Steam in an effort to cripple the Steam client. The new class of malware is being used by developers to cripple Steam users' games and their In early 2013 I decided to use the same set of features that had been included in the first three versions of this blog. The goal of the initial setup was to create a set of scripts that could
ours	A new study of computer simulations suggests that humans and other primates are capable of adapting to a new world without the help of artificial intelligence. The new study, published in the journal Nature, also found that The Internet is a massive, sprawling, interconnected network of connected devices. It is a huge, interconnected network of computers running programs on their own and interacting with each other through the Internet Protocol, a protocol The first time you download an app, you'll want the latest version. If you're running Windows 7 or Windows 8 on a laptop or tablet, you'll need to run it from the same USB The latest version of Windows 10 includes a set of new features for Windows Phone 8.1, such as support for Windows RT and Windows Phone 8.0. These include: The ability to

Table A14: samples of generations in the sentiment control scenario. The target sentiment here is `positive`.

Method	Generation (sentiment = `positive`)
PPO	The chicken is an extremely tasty, versatile, and versatile vegetable that is great when you want to try new flavors. Ingredients: 4 cups of chicken broth 2 cups of sweet, sour
	The horse is a good horse, it's a great companion and the best companion. It's also the horse that I have ever had and I love it. It's a big horse and he's a great
	The lake is one of the most beautiful and most beautiful lakes in the world - a true beauty and the perfect place to get away. This is also an area with great views of the ocean. The
	The pizza is always good and the flavor is just right for me. The crust is really good and the toppings are all fresh, fresh, fresh. I really enjoyed the pizza the first time I was there
GDC	The chicken was a popular delicacy in the Middle Ages and was also an important food to those who ate it. History Early Middle Ages¡—endoftext—¿I'm so excited about this. The
	The horse's life is on the line. In this special episode of the show, a horse's life is on the line as he takes part in a daring battle between two rival factions: the Red and Blue
	The lake has been known to be a popular destination for locals for decades, with its many attractions and attractions, including a lakefront park, lakefront shopping malls, and a beachfront shopping center.
	The pizza shop is in the heart of the heart of downtown Denver. The new restaurant is located in the city's historic downtown. It was built just a few weeks ago as the new home of the
DExperts	The chickenpox vaccine is the lifesaving treatment for both children and adults who receive chickenpox in their homes or hospitals in parts of Africa. Read More An elegant new Indian restaurant celebrates a unique cuisine and
	The horseman Daniel Auldsson is the first Swedish woman to win Olympic silver medal at the World Championships in Beijing in March, while winning gold in the 800m individual medley in Berlin. Australian
	The lake of tears streaming down her cheeks as she listened to her sisters struggle for their lives. An emotional smile dawned across the man's face as he held his daughter in his arms and gently held
	The pizza chain opened in the German city of Hamburg in 2009 with three co-founders – Jens Bohman and Peter Stütte. Their mission is to continue providing exceptional pizzas and delicious food
GeDi	The chicken cresspies has long been associated with chocolate and coffee flavored drinks such as espresso or tea because of their preference for many foods, as well as their tea abilities, eating as much tea as they
	The horse that romped with the big round of sand games which followed took first place. Her Majesty Queen Anne was crowned Queen MP for Newark and it was an inspiration to see her smiling again as
	The lakefront district was formed around Hampden Street, where 18 Chipewyan restaurants cater to Station Market Square, where residents, merchants, restaurateurs, farmers, musicians, entrepreneurs, farmers — often immigrants,
	The pizza chief of the mini-table club, Katara Barella, praised the tailor as one of the most cordial, affectionate and generous people he knows. The pair have lived together on
CTRL	The chicken noodle soup is great, I have tried at least 10 different varieties, I really like this variety. Rating: 5.0 This product was fresh and of high quality. The price in my opinion would be about the same
	The horse shoe style fit on my feet perfectly, no problems at all. After wearing them I did not want to take them off. Rating: 5.0 These are great. Not only do these give great support but they
	The lake that I live on is one of the largest in my state. And at 5,000,000 acres! It gets so deep all you could need and more! After searching for days, including googling "best
	The pizza is so quick and easy, not to mention delicious; plus you can control everything from heat on up. Pizza comes out crisp and yummy. Great for the whole family, even a snack at work. Rating: 5.0
PPLM	The chicken's got the heart! A heart is a heart is a heart. It is a soul! The most beautiful and powerful energy, has the highest potential for the highest, and the
	The horse's great endurance is its ability to endure great pain. A horse's endurance is the ability to endure the greatest pain, which is what makes horses so amazing. Horses have been around long
	The lake in the image is the main weapon in the game. A very good weapon. I like it. I love the way it looks and the way it performs. It is the way it feels and the
	The pizza delivery service is the only way to get a great pizza. We deliver a wide range of delicious pizza from all over the world. Our friendly, helpful staff will make your experience a success for all!
ours	The chicken wings of Indian cuisine are made with fresh spices and spices like coriander, rosehip, rose petal, turmeric, turmeric seeds and spices, together with herbs and spices. Indian cooking
	The horse race will open up to you. You'll get a chance to win tickets, win race tickets, go up to the top of the Grandstand for the race you're passionate for. Plus, we
	The lakefront homes will be developed along the existing land along Grand Avenue, making the home stand out from the larger home. Developers are confident that the existing home will attract potential tenants and will increase
	The pizza place will open on Sunday, September 29. Please bring something special to share with you! You'll need something to bring home to enjoy, plus you'll be able to share with your friends right on

Table A15: samples of generations in the sentiment control scenario. The target sentiment here is `negative`.

Method	Generation (sentiment = **`negative`**)
PPO	The chicken was cooked to a crisp and the sauce was thick. I had to make the noodles for the chicken because they were too thick for them. The sauce was not hot enough. My wife had to make
	The horse is a small, medium, and slow dog. If you have never seen a horse before, the horse is not particularly useful for running because it is slow and has a poor ability to run.
	The lake that the lake was built on, was polluted, and it has a very bad smell. This is the first time there has been an outbreak of this illness, and there has also been an outbreak of
	The pizza was pretty good! But I was really sick of the smell. I had to buy some more. Not sure if this was due to the smell, but the pizza was not tasty. I did not
GDC	The chickenpox virus is spread through infected animals by a single tick, which can cause an infection known as chikungunya. Infections like chikungunya are a major cause
	The horse that killed a woman and injured a man, and that was not an accident, is a real-life horse that's been bred by people who don't want to kill their horses. The
	The lake is a bit of a mess. We were going to go down and clean it up, but I don't really see why we need to. It's pretty much a dead area, and there's
	The pizza delivery man is accused of stealing his girlfriend's $2,500 check and then making her go on a date before he left. According to court documents, the pizza delivery man, who has
DExperts	The chicken died after being cooked. The whole thing is so bad I made chicken breasts from them and then burned them when they were cold. I think I just gave it to a homeless person
	The horse died. The only thing left on my horse was a bunch of rotting corpse shit and a pile of shit on my horse. This mess of bad science. I closed the
	The lake was drained about a quarter mile and half an hour after opening. The problem is there is no lake. "You just ran out the clock." - Jay Leavitt, Lake
	The pizza-eating pig tried to shove his head in. "Why did I do this? Why am I here?!" Unfortunately for this ugly pig, this time around he just looked like a giant pig
GeDi	The chicken market Last month, a Breitbostburg-based skinned worker set herself on fire in an attempt to raise money for "an Orthodox women's charity". The 20-
	The horse crisis spilled over into Utah on Thursday night after local news station KTVK reported on rumors that some UT students were crawling on their lunchouts. "The horse craze ate a whole hog
	The lake looked deserted early Tuesday morning as dozens of hunters headed back to their helicopter while counting on increased water temperatures to cool them off. Some investors sold the trapped-boat assets before sunrise to huddle
	The pizza mess that is the Rohingya is getting worse, not better. But not by much. The Washington Post reports that an attack in a village in northern Myanmar on Saturday resulted in five people
CTRL	The chicken recipe was way over done. I know how to make some good tasting meat sause so not buying this book now. Rating: 1.0 In addition to the lack of a table of contents, you have no
	The horse part of this book isn't bad but the rest is, not so much. Rating: 1.0 I was looking for a cute little western as an intro to equine care and reading my library had many good choices
	The lake was very dirty, there were little fish in the water, but you can see that it is now clean and clear as a glass! You cant even tell how dirty this lake was. Rating: 1.0 These
	The pizza was good. I can tell that it is fresh (should have waited an hour for a few things). But as the other person noted, there were at least 7-8 slices missing from each pizza. They are all
PPLM	The chickenpox epidemic of 1918 in China was the worst in recorded history. In a little over two weeks, the disease killed up to 50 million people in China alone. It also killed off up to 60%
	The horse is the most intelligent of all the mammals. It has a great brain. The animal's sense of smell, which can be so difficult for humans to understand, is a male pteranodon that
	The lake is a place of worship of the Devil, and the Devil, who has the world in his hands. (A. V.) And they were all in a trance, and there was a great light
	The pizza cutter is the world's most famous and most popular knife. Its patented, high-performance, ultra-lightweight, and very cheap, design makes it a great knife to start with. It has
ours	The chickenpox virus is transmitted from infected people to sick or ill persons who have a rash. The disease is usually fatal when people infected with it have a rash. The rash often fades and it may not disappear
	The horse's death is the most tragic accident of the last decade, and it was only the accident that prompted him to leave this horse in his barn. In a statement, Dr. Gaudina
	The lake, which is about half an hour drive from Laguna Niguel, is one of the most polluted parts of the South American country. In 2012, there were about 1,500 dead and
	The pizza and cheeseburger chain, in particular, has been criticized for its low prices. In 2013, a report in the Journal of Market Economics suggested that low-margin pizzas were costing a third fewer