

Why Exposure Bias Matters: An Imitation Learning Perspective of Error Accumulation in Language Generation

Kushal Arora^{1*} Layla El Asri² Hareesh Bahuleyan^{3†} Jackie Chi Kit Cheung^{1‡}

¹Mila / McGill University ²Borealis AI ³Zalando SE

{kushal.arora@mail, jcheung@cs}.mcgill.ca

layla.elasri@borealisai.com, hareeshbahuleyan@gmail.com

Abstract

Current language generation models suffer from issues such as repetition, incoherence, and hallucinations. An often-repeated hypothesis for this brittleness of generation models is that it is caused by the training and the generation procedure mismatch, also referred to as exposure bias. In this paper, we verify this hypothesis by analyzing exposure bias from an imitation learning perspective. We show that exposure bias leads to an accumulation of errors during generation, analyze why perplexity fails to capture this accumulation of errors, and empirically show that this accumulation results in poor generation quality.¹

1 Introduction

Large-scale neural language models have made great strides in a series of language generation tasks such as machine translation (Bahdanau et al., 2014; Vaswani et al., 2017; Raffel et al.), text summarization (See et al., 2017; Lewis et al., 2019; Zhang et al., 2019a), conversational dialog generation (Serban et al., 2015; Lowe et al., 2017; Roller et al., 2020; Zhang et al., 2020), etc.

However, despite the successes achieved by these models on several conditional generation tasks, they continue to suffer from degenerate behaviors such as repetition, a lack of diversity, dullness, and, incoherence, especially in open-ended generation settings such as text completion and dialog modeling (Holtzman et al., 2019). This degenerate behavior is often attributed to a mismatch between the *maximum likelihood training* and gen-

eration procedure mismatch (Welleck et al., 2019; Choi et al., 2020; Li et al., 2016).

Maximum likelihood training also referred to as *teacher forcing* (Williams and Zipser, 1989), factorizes the language model as a linear chain, and maximizes the log-likelihood of this factorized language model on a training corpus. During this maximum likelihood training, the model learns a distribution of the next tokens conditioned on the contexts from the ground-truth training data.

A concern with the MLE-based training is that the ground-truth contexts from the training corpus are not available during generation. Rather, the conditioning contexts during this phase comprise of tokens previously generated by the model itself. The distribution of these contexts seen during the generation phase might be very different from the ones encountered during the training phase. This mismatch is referred to as *exposure bias* (Ranzato et al., 2016; Bengio et al., 2015).

A side-effect of exposure bias is that an error at any step during generation might have a cascading effect as the next context will incorporate this erroneous prediction, deviating away from the ground truth context distribution leading to more errors. These errors will result in sequences that degenerate over the sequence length resulting in incoherent text, lack of vocabulary diversity, and detachment from the source sequence resulting in hallucination, and/or word- and phrase-level repetition.

There is an active debate in the language generation community on the impact of exposure bias in language generation. Authors have both validated (Xu et al., 2019; Zhang et al., 2019b) and questioned (He et al., 2019) the impact of exposure bias on language generation. Several approaches have been proposed to mitigate exposure bias (Ranzato et al., 2016; Shen et al., 2016; Bahdanau et al., 2017; Chen et al., 2020; Leblond et al., 2018; Welleck et al., 2019) but these have neither formalized exposure bias clearly nor provide any empiri-

* A part of this work was done when the author was an intern at Borealis AI.

† During a part of this work, the author was an employee at Borealis AI.

‡ During a part of this work, the author was an Academic Advisor at Borealis AI.

¹Source code to reproduce these experiments is available at https://github.com/kushalarora/quantifying_exposure_bias.

cal evidence that these methods mitigate the effect of exposure bias. Finally, previous works have linked exposure bias to out-of-domain (Wang and Sennrich, 2020) and out of distribution (Schmidt, 2019) generalization, and hallucinations (Wang and Sennrich, 2020) but these claims remain weak in absence of a clear and principled formalization of the exposure bias issue.

In this paper, we attempt to clarify this confusion by formalizing exposure bias in the terms of accumulation of errors and analyzing its impact on generation quality. We do this by providing a theoretically-grounded understanding of the exposure bias issue by analyzing it from an imitation learning perspective. We use this perspective to show that behavior cloning—an imitation learning algorithm is equivalent to teacher forcing under the choice of a particular loss function. We then exploit this equivalence to borrow the bound on error accumulation caused by behavior cloning and use it to quantify exposure bias and analyze the error accumulation in language generation.

Finally, we use this quantifiable definition of exposure bias to demonstrate that models trained using teacher forcing do suffer from an accumulation of errors. We also show, both analytically and empirically, why perplexity fails to capture this error accumulation, and how a lower exposure bias correlates with better generation quality.

2 Language Generation Formulation

Given a finite-sized vocabulary set \mathcal{V} , language generation is posed as a problem of generation a variable-length sequence $w_0^n \in \mathcal{V}^n$ from a language model p_θ , either unconditionally or conditioned on the a source \mathbf{x} , using a decoding algorithm \mathcal{F} .

$$w_0^n = \mathcal{F}(p_\theta; \mathbf{x}) \quad (1)$$

Language modeling is the problem of learning this parameterized model p_θ , that approximates an oracle model o , such that decoding from the model p_θ mimics greedily sampling from the oracle o .

Maximum likelihood-based training factorizes the probability distribution model, $p_\theta(w_0^n)$, into a linear chain, i.e.,

$$p_\theta(w_0^n; \mathbf{x}) = \prod_{i=1}^n p_\theta(w_i | w_0^{i-1}; \mathbf{x}) p(w_0), \quad (2)$$

where w_i is the token to be generated at step i and w_0^{i-1} is the context at time i ; i.e., all the tokens

seen from step 0 to step $i - 1$.²

During maximum likelihood training, the language model is trained by minimizing the negative log-likelihood on the corpus \mathcal{D} , i.e.,

$$\theta^* = \operatorname{argmin}_\theta \frac{-1}{|\mathcal{D}|} \sum_{w_0^n \in \mathcal{D}} \sum_{i=0}^n \log p_\theta(w_i | w_0^{i-1}), \quad (3)$$

where $|\mathcal{D}|$ is the number of tokens in the corpus.

Given a trained language model p_θ , the simplest strategy for generating a target sequence is to greedily sample the model i.e., at each step i , pick the most probable token $w_i = \operatorname{argmax} p_\theta(\cdot | w_0^{i-1}; \mathbf{x})$ —as its prediction. For the next step $i + 1$, we use w_i to generate the context $w_0^i = w_0^{i-1} w_i$, and use it to predict the next token. This continues either until the maximum sequence length (T) is reached, or a special end of sequence token (EOS) is generated.

3 An Imitation Learning Perspective of Language Generation

In this section, we will present an imitation learning perspective of language generation. This framing will allow us to borrow theoretical machinery from imitation learning literature to formalize the exposure bias issue and quantify it in terms of accumulation of error due to procedural mismatch between MLE-based training and generation.

We start by posing language generation as a sequential decision-making problem and language modeling as an instance of imitation learning. We exploit these parallels to show behavior cloning—an imitation learning algorithm, is equivalent to teacher forcing under a choice of a particular loss function. We then exploit this equivalence to quantify the error accumulation due to exposure bias.

Language Generation is a Sequential Decision-Making Problem:

A sequential decision-making problem can be formalized as learning a policy $\pi(a_t | s_t)$ over a space of actions $a_t \in \mathcal{A}$ and states $s_t \in \mathcal{S}$ where the next state s_{t+1} is conditioned on the current state-action pair and is determined by the transition distribution $P(s_{t+1} | s_t, a_t)$. We can use this framework to pose language generation as an instance of a sequential decision-making problem with language model p_θ as the policy, contexts $w_0^{t-1} \in \mathcal{V}^*$ as states, the next token prediction $w_t \in \mathcal{V}$ as actions, and concatenation function as the transition function.

²As w_0 is usually a fixed SOS token, $p(w_0) = 1$. We will drop $p(w_0)$ from the subsequent equations for brevity.

This perspective allows us to appreciate the fact that, during generation, predictions at previous steps affect the next predictions, and error over time can cascade resulting in incoherent sequences.

Language Modeling is Imitation Learning: Imitation learning is a class of methods to solve a sequential decision-making problem while having access to the oracle policy o or data generated by the oracle, i.e., $\mathcal{D} = \{(s_t, a_t) | s_t \sim d_o^t, a_t \sim o(\cdot | s_t)\}$. Here, d_o^t is the oracle induced state-visitation distribution at time t . In imitation learning, an agent learns a model policy π that reproduces the expert policy o but on the state-visitation distribution d_π^t that has been induced by the model policy π .

The sequential decision-making perspective of language generation allows us to pose language modeling as an instance of imitation learning—learning a model for a sequential decision-making problem with the help of an expert oracle (in RL-based methods) or using the data generated by the oracle (for MLE-based methods).

Teacher Forcing is Behavior Cloning: The assumption of access to an oracle is unrealistic in many scenarios. Behavior cloning is an approach to solve an imitation learning problem using only the training data generated by an oracle. In this setup, the state-action pairs in the training data are assumed to be identically and independently distributed. This is equivalent to reducing a sequential decision-making problem to a supervised multi-class classification learning problem.

Concretely, this learning problem can be seen as minimizing the expected per-step loss under the state distribution induced by the oracle:

$$L^{BC}(\pi) = \sum_{t=1}^T \mathbb{E}_{s_t \sim d_o^t} [l(s_t, \pi; o)] \quad (4)$$

$$\approx \frac{-1}{|\mathcal{D}|} \sum_{(s_t, a_t) \in \mathcal{D}} l(s_t, \pi; o), \quad (5)$$

Here, $L^{BC}(\pi)$ is the behavior cloning loss and $l(s, \pi; o)$ is the per-step loss.

Similarly, in practical scenarios, language models are also trained on a finite training corpus, \mathcal{D} , that assumed to be generated by the oracle, i.e., $\mathcal{D} = \{(w_0^{t-1}, w_t) | w_0^{t-1} \sim d_o^{t-1}, w_t \sim o(\cdot | w_0^{t-1})\}$.

Maximum likelihood training loss from Equation 3, can be reformulated as learning the distribution over the next tokens, conditioned on the train-

ing contexts generated by the oracle, $w_0^{t-1} \sim d_o^{t-1}$:

$$L^{TF}(p_\theta) = \frac{-1}{|\mathcal{D}|} \sum_{(w_0^{i-1}, w_i) \in \mathcal{D}} \log p_\theta(w_i | w_0^{i-1}), \quad (6)$$

$$\approx \sum_{t=1}^T \mathbb{E}_{w_0^{t-1} \sim d_o^t} [-\log p_\theta(w_i | w_0^{i-1})] \quad (7)$$

The behavior cloning loss (Equation 4) is equivalent to the language modeling loss (Equation 7) with $l(p_\theta, w_0^{t-1}; o) = -\log p_\theta(w_i | w_0^{i-1})$.

For our analysis though, we define per-step loss for language modeling, $l(p_\theta, w_0^{t-1}; o)$ as:

$$l(p_\theta, w_0^{t-1}; o) = \mathbb{E}_{w_t \sim o(\cdot | w_0^{t-1})} \log \frac{o(w_t | w_0^{t-1})}{p_\theta(w_t | w_0^{t-1})}, \quad (8)$$

This definition ensures that the per-step loss for oracle is zero, i.e., $l(o, w_0^{t-1}; o) = 0$.

The per-step loss function defined by equation 8 ensures that the behavior cloning loss, $L^{BC}(p)$, under our definition is equivalent to teacher forcing loss, $L^{TF}(p)$, up to a constant term. This equivalence of $L^{BC}(p)$ and $L^{TF}(p)$ ensures that the model learned by minimizing either of the two losses will be identical.

Language Generation is Policy Rollouts: During policy rollouts, an agent in state s_t , executes the action a_t , sampled from policy π , and ends up in state s_{t+1} . The agent’s next state is dependent upon its own actions. This state evolution can be formulated as sampling from state-visitation distribution induced by the policy π , i.e., $s_{t+1} \sim d_\pi^{t+1}$.

The performance of policy π during rollouts can be measured using the loss (cost) of executing the policy π , i.e.,

$$L^I(\pi) = \sum_{t=1}^T \mathbb{E}_{s_t \sim d_\pi^t} [l(s_t, \pi; o)] \quad (9)$$

We can also formulate language generation in terms of policy rollouts from imitation learning. Mathematically, we can express generation as sampling contexts from model’s context distribution, i.e., $w_0^{j-1} \sim d_{p_\theta, \mathcal{F}}^j$, and generating the next token w_j conditioned on w_0^{j-1} , using the decoding algorithm \mathcal{F} , i.e.,

$$\{w_j = \mathcal{F}(p_\theta, w_0^{j-1}) | w_0^{j-1} \sim d_{p_\theta, \mathcal{F}}^j\} \quad (10)$$

We can now define the inference-time loss, $L^I(p_\theta)$, as the accumulated T -step loss of model p_θ imitating oracle o on the context distribution induced by the model:

$$L^I(p_\theta) = \sum_{t=1}^T \mathbb{E}_{\substack{w_0^{t-1} \sim d_{p_\theta, \mathcal{F}}^t \\ w_t \sim o(\cdot | w_0^{t-1})}} \log \frac{o(w_t | w_0^{t-1})}{p_\theta(w_t | w_0^{t-1})}, \quad (11)$$

where $d_{p_\theta, \mathcal{F}}^t(w_0^t) := p_\theta(w_0^{t-1})$, is the context distribution at step t , induced due to use of model p_θ and the decoding algorithm \mathcal{F} , from step 1 to $t-1$.

4 Exposure Bias and Error Accumulation

Ranzato et al. (2016) defined exposure bias as a behavioral mismatch between maximum likelihood-based training and generation procedure. During maximum likelihood-based training, the next token distribution is conditioned on ground truth data whereas, during generation, it has to rely on the model's own previously generated tokens. They also postulated that this training and generation context distribution mismatch might result in an accumulation of errors during generation.

Intuitively, when the model produces a token w_i that makes the resulting context w_0^i unfamiliar, it might not be able to continue the generation adequately and is likely to produce another token which will further make the context flawed. This phenomenon reinforces itself as the context drifts further from what the oracle would produce, leading to an accumulation of errors.

In the imitation learning literature, the accumulation of errors while rolling out a policy trained using behavior cloning is analyzed in the terms of inference-time regret of the behavior cloning policy, π_{BC} , with respect to the oracle policy, o , (Ross and Bagnell, 2010; Ross et al., 2011) i.e.,

$$\mathcal{R}(\pi_{BC}) = L^I(\pi_{BC}) - L^I(o) \quad (12)$$

Let ϵ_t be the expected error of executing policy π at step t on the state-visitation distribution induced by the oracle o , i.e.,

$$\epsilon_t = \mathbb{E}_{s \sim d_o^t} [l(s, \pi; o)] \quad (13)$$

Let ϵ be the average expected error of executing policy π over T step, i.e., $\epsilon = 1/T \sum_{t=1}^T \epsilon_t$. Assuming $l(s, \pi, o)$ is an upper bound on $[0, 1]$ loss, we can bound the regret for a policy π_{BC} as,

$$T\epsilon \leq \mathcal{R}(\pi_{BC}) \leq T^2\epsilon. \quad (14)$$

The lower-bound in Equation 14 assumes no accumulation of error, hence an expected error of ϵ at each step, whereas the upper bound assumes the worst-case scenario, resulting in linear growth in error at each step and overall quadratic accumulative growth w.r.t. maximum sequence length T .

Relying on the imitation learning perspective of language generation presented in the previous section, we can now borrow this regret-based analysis from imitation learning literature to similarly bound the regret of a language generation model as

$$T\epsilon \leq \mathcal{R}(p_\theta, \mathcal{F}) \leq T^2\epsilon. \quad (15)$$

where p_θ is the model being used for generation, \mathcal{F} is the decoding method being used for generation, $\epsilon = 1/T \sum_{t=1}^T \epsilon_t$ and ϵ_t is defined as

$$\epsilon_t = \mathbb{E}_{\substack{w_0^{t-1} \sim d_o^t \\ w_t \sim o(\cdot | w_0^{t-1})}} \log \frac{o(w_t | w_0^{t-1})}{p_\theta(w_t | w_0^{t-1})} \quad (16)$$

We will now use these bounds on the regret to analyze and quantify the error accumulation due to exposure bias in language generation.

5 Quantifying Error Accumulation due to Exposure Bias

In our analysis, we use two metrics, $\text{AccErr}_{\leq}(l)$ and $\% \text{ExAccErr}_{\leq}(l)$ to measure the impact of error accumulation due to exposure bias.

We define accumulated errors up to length l , $\text{AccErr}_{\leq}(l)$, as

$$\text{AccErr}_{\leq}(l) = \mathcal{R}_{\leq l}(p_\theta, \mathcal{F}) / \epsilon_{\leq l} \quad (17)$$

Here, $\mathcal{R}_{\leq l}(p_\theta, \mathcal{F})$ be the regret due to the use of language model p_θ and decoding method, \mathcal{F} , up to sequence length l , and $\epsilon_{\leq l} = 1/l \sum_{t=1}^l \epsilon_t$ is the expected per-step error up to length l .

This metric captures the growth of error w.r.t. sequence length l . If exposure bias does indeed leads to error accumulation, $\text{AccErr}_{\leq}(l)$ should grow super-linearly w.r.t. l .

We define our second metric, $\% \text{ExAccErr}_{\leq}(l)$, as percentage of excess errors committed by the model that can be attributed to exposure bias, i.e.,

$$\% \text{ExAccErr}_{\leq}(l) = \frac{\mathcal{R}_{\leq l}(p_\theta, \mathcal{F}) - l\epsilon_{\leq l}}{l\epsilon_{\leq l}} * 100$$

Here, $l\epsilon_{\leq l}$ is the lower bound on the regret and is the minimum number of errors (ϵ per step) a model would make if there was no accumulation of errors.

$\%ExAccErr_{\leq}(l)$ allows us to compare models, training algorithms, and decoding strategies on the extra error that might be caused/mitigated by their use. A model, training algorithm, or a decoding strategy that perfectly mitigates the exposure bias will result in zero excess accumulated errors.

In the rest of the paper, we use these definitions to show: 1.) error accumulation in language generation is real, 2.) perplexity fails to capture this error accumulation, 3.) lower exposure bias correlates with a higher quality generation that is more coherent, uses diverse vocabulary, and is less repetitive.

6 Study Setup: Open-ended Generation

Text Completion Setup: Text completion is the standard experimental setup to measure the quality of generation in open-ended language generation (Holtzman et al., 2019; Welleck et al., 2019). It is also a generalization of a numerous practical language generation applications such as story generation (Fan et al., 2018), contextual text completion (Radford et al., 2019), dialog modeling (Zhang et al., 2018), etc.

Text completion models take a text passage or prefix $w_0^j \sim o$ as an input and generate a coherent continuation of the prefix, w_{j+1}^n using the language model p_θ and the decoding algorithm \mathcal{F} , i.e., $w_{j+1}^n = \mathcal{F}(p_\theta, w_0^j)$. In this paper, we use this text-completion setup to analyze the error accumulation due to exposure bias and its correlation with language generation quality.

Language Model and Dataset: We conduct our analysis using the GPT2 language model (Radford et al., 2019). We use the GPT2-117M model as our evaluation language model and use train split of Wikitext-103 (Merity et al., 2016) for prompts. We rely on GPT-2 model fine-tuned on Wikitext-103 as our approximate oracle. We tokenize the Wikitext-103 dataset using GPT-2’s tokenization scheme. We chunk Wikitext-103’s train split into sequences of length 512. Of these, we use the first 50 tokens as prompts for our generation experiments and generate the completions to a maximum length of 512 or up to the end of the sequence token. We use a total of $20k$ prompts for our evaluation.

7 Results

7.1 Error Accumulation in Language Generation is Real!

Figure 1a plots $AccErr_{\leq}(l)$ w.r.t. sequence length, l . The support (dotted, orange line) $y = x$, captures the linear growth. It shows $AccErr_{\leq}(l)$ grows near-quadratically w.r.t. sequence length, empirically validating the theory that exposure bias would lead to accumulation of errors. Figure 1b, further strengthens this claim by demonstrating near-linear growth in excess errors w.r.t. to sequence length.

We hypothesize that these excess errors would manifest in the form of language degeneration, especially in the latter part of the sequence, and would cause issues such as hallucinations, limited vocabulary, and word- and phrase-level repetitions.

7.2 Perplexity is Not Enough

Perplexity is a standard measure used to evaluate the quality of a language model. It is often used as a proxy measure for the text generation quality of the language model. In this section, we argue perplexity paints an incomplete picture regarding a model’s ability to generate high-quality coherent text. It only captures the average per-step error generalization gap (or lack of it) but fails to account for the error accumulation due to exposure bias. These accumulated errors, as seen in the previous section, can grow near-quadratically and can prove to be a major concern for any generation model that generates sequences longer than a few words.

Perplexity can be seen as scaled exponentiated average per-step error, ϵ , computed over a held-out test set, \mathcal{D}_h :

$$\epsilon = 1/T \sum_{t=1}^T \mathbb{E}_{\substack{w_0^{t-1} \sim d_o \\ w_t \sim o(\cdot|w_0^{t-1})}} \log \frac{o(w_t|w_0^{t-1})}{p(w_t|w_0^{t-1})}. \quad (18)$$

$$\approx \frac{-1}{|\mathcal{D}_h|} \sum_{(w_0^{i-1}, w_i) \in \mathcal{D}_h} \log p_\theta(w_i|w_0^{i-1}) + c, \quad (19)$$

$$= H(p_\theta; \mathcal{D}_h) + c. \quad (20)$$

where $H(p_\theta; \mathcal{D}_h)$ is the entropy rate (log perplexity) of the model p_θ on held-out test set \mathcal{D}_h .

As entropy rate is a linear function of average per-step error, we hypothesize that it will only be able to measure the per-step generalization gap of the model and will fail to capture the error accumulation caused by reducing a sequential decision-making problem to a supervised learning problem.

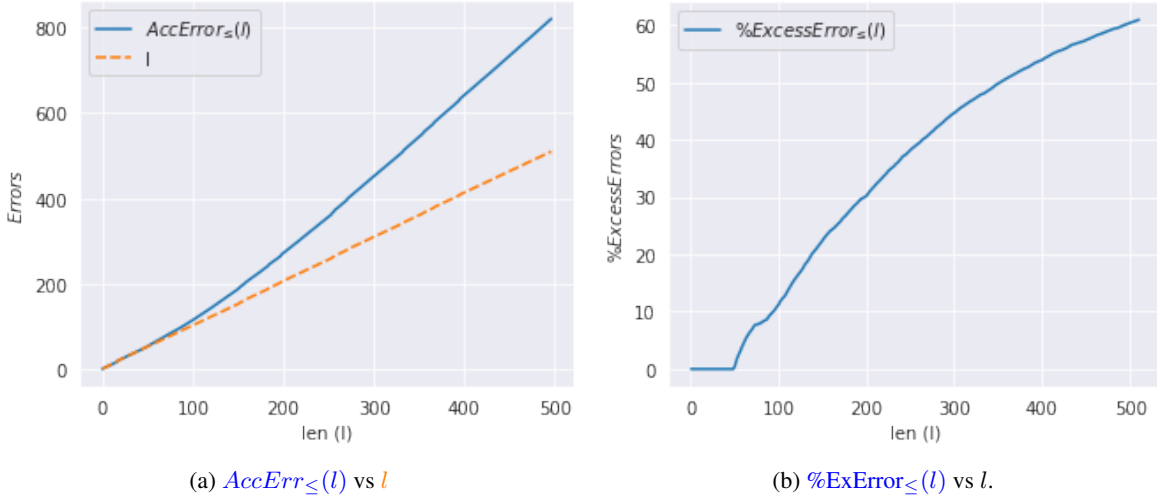


Figure 1: Figure 1a plots accumulated error till len l ($AccErr_{\le}(l)$) w.r.t. l . This graph shows the quadratic growth of accumulated errors w.r.t to sequence length (l) as predicted by the theory. Figure 1b plots % excess errors due to error accumulation ($\%ExError_{\le}(l)$) caused by exposure bias. This indicates that **extra errors** due to exposure bias grows near-linearly with the sequence length, and decoding using greedy search results in over 70% more errors.

Search	Generation Quality				
	%ExErrAcc (\downarrow)	seq-rep-4 (\downarrow)	rep (\downarrow)	wrep (\downarrow)	uniq (\uparrow)
Greedy	60.96%	0.8990	0.4423	0.4136	7833
Beam (k=5)	69.72%	0.8094	0.4064	0.3787	10966

Sampling					
w/ Temp (temp=1)	39.37%	0.1883	0.2547	0.2301	23729
w/ Temp (temp=1.2)	24.75%	0.1556	0.2271	0.2033	25225
w/ top-k (k=100)	35.37%	0.1690	0.2409	0.2166	26251
w/ top-p (p=0.94)	48.71%	0.2218	0.2743	0.2490	22582
Human	-	0.0274	0.4338	-	28739

Table 1: Impact of error accumulation on generation quality. We observe that stochastic decoding methods not only lead to diverse language generation but also have lower exposure bias than the deterministic methods.

In Figure 2, we plot the entropy rate, $H(p_{\theta}; \mathcal{D}_h)_{\leq l}$, w.r.t. average per-step error, $\epsilon_{\leq l}$, and length-normalized regret up to length l , $\mathcal{R}_{\leq l}(p_{\theta}, \mathcal{F})/l$. We observe a strong correlation between the entropy rate and average per-step error ($\rho = 0.9997$) validating our theoretical observation that perplexity can capture the per-step generalization gap of language model p_{θ} . On the other hand, the length-normalized regret exhibit poor correlation with the entropy rate ($\rho = 0.4003$) indicating perplexity’s failure to capture the error accumulation due to exposure bias.

A case in point of perplexity’s inability to capture error accumulation is the degenerate behavior of GPT-2 (Radford et al., 2019) while generating moderately long sequence under greedy or beam

search. This happens despite GPT2 having a low zero-shot perplexity on the held-out set of Wikitext-103 dataset (perplexity: 37.50), but despite this, it suffers from degeneration issues such as repetition, low vocabulary usage, and a lack of coherent generation. We hypothesize that the degenerate behavior of large pre-trained language models such as GPT-2 under greedy or beam search is the result of this accumulation of errors. An example of this behavior is presented in Table 2 where we observe GPT2 generating repetitive and incoherent text completion for a WikiText-103 prompt under deterministic decoding schemes such as greedy and beam decoding.




Figure 2: Analyzing (log) perplexity ($H_{<l}$) w.r.t to average per-step error ($\epsilon_{<l}$), and length-normalized exposure bias regret ($\mathcal{R}_{<l}(p_\theta, \mathcal{F})/l$). We observe that perplexity strongly correlates with average per-step error ($\rho = 0.9997$), but it has a weaker correlation with length-normalized regret ($\rho = 0.4003$).

7.3 Error Accumulation impacts Generation Quality

Finally, we examine the hypothesis that poor text generation capabilities of pre-trained large language models under greedy decoding might be due to the error accumulation caused by a procedural mismatch between generation and maximum likelihood training (Vijayakumar et al., 2016; Welleck et al., 2019; Holtzman et al., 2019).

The regret-based definition of error accumulation allows us to analyze exposure bias along two axes of variation: the trained language model, p_θ , and the decoding algorithm, \mathcal{F} . In this set of experiments, we explore the impact of various decoding schemes on error accumulation due to exposure bias and the quality of the completed text.

Benchmarking various decoding strategies allows us to verify if an accumulation of errors does indeed lead to degeneration as the choice of decoding algorithm does not impact the average per-step error, ϵ , or held-out test set perplexity ($H(p_\theta; \mathcal{D}_h)_{<l}$). This rules out the role of modeling and model training in language degeneration across different decoding algorithms. Hence, it is reasonable to causally link a decoding algorithm’s resultant improvement on language generation to its ability to reduce error accumulation.

For quantitative analysis of the impact of various decoding algorithms on quality of language generation, we measure the completion quality by using

Context: ... He made one attempt to break through the blockade on 23 December, but it proved impossible to do as the Union controlled every channel from Savannah, aided ...

Greedy: ... Savannah, aided by the British, to the North. The blockade was broken on 23 December by the British and the Americans. **The blockade was broken on 23 December by the British and the Americans. The blockade was broken on 23 December by the British and the Americans. The blockade was broken on 23 December by the British and the ...**

BeamSearch (k=5): ... Savannah, aided by the North Carolina Railroad. In the spring of 1788, in the midst of the blockade, he made a move to join the Confederacy, **but it was too late.** On 23 December 1788, he marched to the battle of Custer’s Creek in the South Carolina River. In the following weeks, he returned to Savannah **on 23 December, but it was too late.** On 24 December 1788, he made another attempt to break through the blockade **on 23 December, but it was too late.** ...

Sampling w/ Temp. (temp=1.2): ... Savannah, aided by US Navy ships, and took him by surprise. One of the most famous incidents involved General Grant’s arrest at the Battle of Camp Amish on 16 August. He was being held by the Union Army in Atlanta and accused of taking down General Joseph E. Douglas, a Marine commander who was in command of Fort Hood, an Army outpost near Fort Laramie, Mississippi. ...

Top-k Sampling (k=100): ... Savannah, aided by a little naval warplane bombardment which had crippled many parts of the ship. As an early American merchant vessel he was carried forward as a heavy merchant, while at some time it was necessary to take up residence in Savannah to be used as a regular transport vessel. In the meantime, on the 20th of January 1816, Jefferson died. He was buried in an unmarked grave in the chapel inside Savannah in the family home at Glenbrook, where the rest of the family resides. ...

Top-p Sampling (p=0.94): ... Savannah, aided in the capture of his land and his supplies. He made two attempts to stop a blockade on 28 December, but failed to make it. There were, however, six people in Savannah who were convinced that Captain Walker was the cause of the blockade. He then made a trip to North Carolina where he gave up hope. ...

Gold: ... He made one attempt to break through the blockade on 23 December, but it proved impossible to do as the Union controlled every channel from Savannah, aided by their occupation of Tybee Island at the mouth of the Savannah River. Bulloch reported to Mallory in late January 1862 that breaking out was hopeless so Mallory ordered him to turn the ship over to another officer and to return to Europe some other way. ...

Table 2: Examples of completions using various decoding methods. We observe that the deterministic decoding schemes produce less diverse, incoherent, and more repetitive (highlighted in red) text.

the same metrics as Welleck et al. (2019), and these metrics are: 1.) **rep/128** measures if the prediction token at step t occurs in previous 128 steps, 2.) **wrep/128** counts the prediction’s repetition at step t only if the predicted token is not the ground-truth token at that position, 3.) **seq-rep-4** measure the repetition at the 4-gram level, and 4.) **uniq** measure the vocabulary diversity by accounting for the number of unique tokens generated by the model.

Table 1 shows that the use of stochastic sampling-based decoding algorithms results in diverse and more coherent language generation and a lower percentage of excess errors. Sampling with temperature (with temp=1.2) leads to the least amount of repetition (both at the token and the n-gram level), second highest vocabulary diversity, and the least amount of excess errors due to exposure bias. This also bears out from our qualitative analysis in Table 2 as sampling with temperature produces the most coherent text. Deterministic decoding schemes, in contrast, fare poorly in both reducing exposure bias and on language generation quality metrics, producing repetitive and incoherent text. These quantitative and qualitative experiments offer us evidence that reducing exposure bias does lead to more coherent text generation.

We hypothesize that the reasonable amount of randomness introduced by stochastic sampling helps the model avoid sampling the most likely token at each time step, thus avoiding possible divergent contexts that might have resulted in a degenerate completion in the future. We conjecture that this timely intervention restricts the generation context distribution from moving too far away from training context distribution helping it avoid compounding of errors. This is also borne out by qualitative analysis as a reasonable amount of stochasticity does result in texts which look more coherent and oracle-like. A broader analysis of this behavior though is beyond the scope of this work and is left for future work.

8 Related Work

Non-MLE Training Methods: Several approaches have been proposed to mitigate the exposure bias issue including RL-based optimization objectives (Ranzato et al., 2016; Shen et al., 2016; Bahdanau et al., 2017; Chen et al., 2020), learning to search (Leblond et al., 2018), energy-based models (Deng et al., 2020), imitation learning (Du and Ji, 2019), generative adversarial networks (Yu et al., 2017) and knowledge distillation (Liu et al., 2019). Although these methods motivate their approaches intending to reduce exposure bias, they neither formalize exposure bias clearly nor provide any empirical evidence that these methods mitigate the effect of exposure bias. In this paper, we proposed a quantifiable definition of exposure bias by analyzing the issue from a principled imitation learning perspective. This definition can be used or

adapted to evaluate the various novel training and modeling approaches on their ability to do reduce exposure bias.

Smarter Decoding Methods: Large language models have unusually low test perplexities but they falter at coherent and diverse language generation specifically in open-ended language generation tasks especially while using deterministic decoding schemes. Several authors (Vijayakumar et al., 2016; Welleck et al., 2019; Holtzman et al., 2019) have hypothesized that training and inference mismatch due to MLE-based training is responsible for the degenerate behavior. They have proposed smarter decoding schemes to mitigate the side effects of exposure bias resulting in better generation quality. Despite this being an active area of research, this often-repeated hypothesis for degenerate generation behavior has not received serious treatment till now. In this paper, we take a step towards explaining this discrepancy and show that error accumulation due to exposure bias might be the reason for this degenerate behavior and explain why perplexity has a handicap in capturing this compounding of errors.

Analyzing Exposure Bias: Schmidt (2019) and Wang and Sennrich (2020) link exposure bias to generalization gap due to distribution and domain shift respectively. Performance degradation under domain and distribution shift is a major issue with language generation, and direct evidence supporting this hypothesis will provide insights into building more robust language generation models. Unfortunately, neither of the papers formalize the notion of exposure bias or empirically link the generalization gap to exposure bias directly.

Three recent papers, Xu et al. (2019); Zhang et al. (2019b); He et al. (2019), have tried to empirically evaluate the impact of exposure bias on language generation. The first two papers validate the existence of exposure bias whereas He et al. (2019) show language models have self-recovering ability negating the impact of exposure bias. All three analyses are based on the empirical definition of exposure bias which, in turn, is based on the informal formulation by Ranzato et al. (2016).

In this paper, we provide a principled and theoretically grounded approach to analyze exposure bias in language generation and show that it is indeed a problem and that it might explain the degeneration issue with large language models on open-ended tasks under deterministic decoding.

9 Discussion

In this paper, we analyze language generation from an imitation learning perspective. We use this analysis to arrive at a theoretical bound on error accumulation due to exposure bias. This bound predicts a super-linear growth in errors accumulation during generation due to exposure bias. In our experiments, we validate this bound and show that accumulation due to exposure bias indeed results in super-linear growth in errors.

We then show, both analytically and empirically, why perplexity is not enough to capture this accumulation of errors and hypothesize that this accumulation of errors is responsible for the degenerate language generation. Finally, we provide some evidence for this hypothesis by evaluating the impact of various decoding schemes on error accumulation and generation quality. We show that techniques that improve the generation quality do result in a lower error accumulation and causally link language generation quality to error accumulator due to exposure bias.

Our analysis provides a principled and theoretically grounded way to understand the exposure bias. We believe this analysis can pave way for developing smarter training and decoding algorithms to address this error accumulation resulting in more robust language generation models.

Acknowledgments

We would like to thank the reviewers for their valuable feedback. This work is supported by funding from Samsung Electronics. The last author is supported by the Canada CIFAR AI Chair program. This research was enabled in part by support provided by Calcul Québec³, and Compute Canada⁴. We would also like to thank Khimya Khetarpal, Sachin Grover, Ankit Anand, Jayakumar Subramanian for feedback on current and previous drafts of this paper, and colleagues at Borealis AI for their valuable inputs and discussions during the first author’s internship at Borealis AI.

References

Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. *An actor-critic*

³<https://www.calculquebec.ca>

⁴<https://www.compute-canada.ca>

algorithm for sequence prediction. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. *Neural Machine Translation by Jointly Learning to Align and Translate*. *arXiv:1409.0473 [cs, stat]*. ArXiv: 1409.0473.

Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. *Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks*. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1171–1179. Curran Associates, Inc.

Yu Chen, Lingfei Wu, and Mohammed J. Zaki. 2020. *Reinforcement learning based graph-to-sequence model for natural question generation*. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.

Byung-Ju Choi, Jimin Hong, David Park, and Sang Wan Lee. 2020. *F²-Softmax: Diversifying Neural Text Generation via Frequency Factorized Softmax*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9167–9182, Online. Association for Computational Linguistics.

Yuntian Deng, Anton Bakhtin, Myle Ott, Arthur Szlam, and Marc Aurelio Ranzato. 2020. *Residual energy-based models for text generation*. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.

Wanyu Du and Yangfeng Ji. 2019. *An empirical comparison on imitation learning and reinforcement learning for paraphrase generation*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 6011–6017. Association for Computational Linguistics.

Angela Fan, Mike Lewis, and Yann Dauphin. 2018. *Hierarchical Neural Story Generation*. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.

Tianxing He, Jingzhao Zhang, Zhiming Zhou, and James R. Glass. 2019. *Quantifying exposure bias for neural language generation*. *CoRR*, abs/1905.10617.

Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2019. *The Curious Case of Neural Text Degeneration*. *arXiv:1904.09751 [cs]*. ArXiv: 1904.09751.

- Rémi Leblond, Jean-Baptiste Alayrac, Anton Osokin, and Simon Lacoste-Julien. 2018. [SEARNN: training rnns with global-local losses](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension](#). *arXiv:1910.13461 [cs, stat]*. 00007 arXiv: 1910.13461.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A Diversity-Promoting Objective Function for Neural Conversation Models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Rui Liu, Berrak Sisman, Jingdong Li, Feilong Bao, Guanglai Gao, and Haizhou Li. 2019. [Teacher-student training for robust tacotron-based TTS](#). *CoRR*, abs/1911.02839.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. [Towards an Automatic Turing Test: Learning to Evaluate Dialogue Responses](#). *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1:1116–1126.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. [Pointer Sentinel Mixture Models](#). *arXiv:1609.07843 [cs]*. 00218 arXiv: 1609.07843.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. page 67.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence level training with recurrent neural networks](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M. Smith, Y.-Lan Boureau, and Jason Weston. 2020. [Recipes for building an open-domain chatbot](#). *arXiv:2004.13637 [cs]*.
- Stéphane Ross and Drew Bagnell. 2010. [Efficient reductions for imitation learning](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, volume 9 of *JMLR Proceedings*, pages 661–668.
- Stephane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. [A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning](#). In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 627–635.
- Florian Schmidt. 2019. [Generalization in Generation: A closer look at Exposure Bias](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 157–167, Hong Kong. Association for Computational Linguistics. 00002.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get To The Point: Summarization with Pointer-Generator Networks](#). *arXiv:1704.04368 [cs]*. ArXiv: 1704.04368.
- Iulian V. Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, and Joelle Pineau. 2015. [Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models](#). *arXiv:1507.04808 [cs]*. ArXiv: 1507.04808.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. [Minimum Risk Training for Neural Machine Translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1683–1692, Berlin, Germany. Association for Computational Linguistics. 00202.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention Is All You Need](#). *arXiv:1706.03762 [cs]*. ArXiv: 1706.03762.
- Ashwin K. Vijayakumar, Michael Cogswell, Ramprasath R. Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. [Diverse Beam Search: Decoding Diverse Solutions from Neural Sequence Models](#).
- Chaojun Wang and Rico Sennrich. 2020. [On exposure bias, hallucination and domain shift in neural machine translation](#). *CoRR*, abs/2005.03642.
- Sean Welleck, Ilia Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. [Neural Text Generation with Unlikelihood Training](#). *arXiv:1908.04319 [cs, stat]*. 00000 arXiv: 1908.04319.
- Ronald J. Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280.

- Yifan Xu, Kening Zhang, Haoyu Dong, Yuezhou Sun, Wenlong Zhao, and Zhuowen Tu. 2019. [Rethinking exposure bias in language modeling](#). *CoRR*, abs/1910.11235.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. [Seqgan: Sequence generative adversarial nets with policy gradient](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 2852–2858. AAAI Press.
- Amy Zhang, Adam Lerer, Sainbayar Sukhbaatar, Rob Fergus, and Arthur Szlam. 2018. [Composable Planning with Attributes](#). *arXiv:1803.00512 [cs]*. ArXiv: 1803.00512.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2019a. [Pegasus: Pre-training with extracted gap-sentences for abstractive summarization](#).
- Wen Zhang, Yang Feng, Fandong Meng, Di You, and Qun Liu. 2019b. [Bridging the gap between training and inference for neural machine translation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4334–4343. Association for Computational Linguistics.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. [DIALOGPT : Large-Scale Generative Pre-training for Conversational Response Generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.