

# Correction with Backtracking Reduces Hallucination in Summarization

Zhenzhen Liu\* Chao Wan Varsha Kishore Jin Peng Zhou

Cornell University

{zl535, cw862, vk352, jz563}@cornell.edu

Minmin Chen

Google DeepMind

minminc@google.com

Kilian Q. Weinberger

Cornell University

kqw4@cornell.edu

## Abstract

Abstractive summarization aims at generating natural language summaries of a source document that are succinct while preserving the important elements. Despite recent advances, neural text summarization models are known to be susceptible to hallucinating (or more correctly confabulating), that is to produce summaries with details that are not grounded in the source document. In this paper, we introduce a simple yet efficient technique, CoBa, to reduce hallucination in abstractive summarization. The approach is based on two steps: hallucination detection and mitigation. We show that the former can be achieved through measuring simple statistics about conditional word probabilities and distance to context words. Further, we demonstrate that straight-forward backtracking is surprisingly effective at mitigation. We thoroughly evaluate the proposed method with prior art on three benchmark datasets for text summarization. The results show that CoBa is effective and efficient in reducing hallucination, and offers great adaptability and flexibility. Code can be found at <https://github.com/zhenzhel/CoBa>.

## 1 Introduction

Recent summarization methods, based on neural sequence-to-sequence and language models (LM), are able to produce high-quality summaries (Zhang et al., 2020; Chung et al., 2022; Touvron et al., 2023a). However, despite their impressive capabilities these summarization models are prone to hallucinations, a phenomenon where models make statements that seem plausible but are not grounded in the source document (Pagnoni et al., 2021a; Maynez et al., 2020a; Zhao et al., 2020). Hallucinations compromise the accuracy and trustworthiness of the generated summaries.

We hypothesize that one reason for hallucination is that sometimes after a LM generates partial text, there is no completion that is grounded in the source text. An illustration of this situation is shown in Figure 1. Although the partial sentence *I live in* is highly plausible, it forces the LM to specify where the person lives, even though this is not specified in the source document. Such situations can often be *detected* by intrinsic properties of hallucinated text: (1) the first word of a hallucinated sequence tends to have low conditional probability, (2) hallucinations are not supported by words in the context, and therefore have a large distance to context words. Returning to our previous example, if the language model continues the sentence *I live in* without any support from the context, *Munich* might be just as plausible as *New York*, or *Penn State*. None of the locations would have particularly high probability, therefore triggering condition (1). Further, if none of the cities are mentioned in the context, all would have a large word distances to the context words, triggering condition (2). Once the beginning of a hallucination is detected, we *backtrack* and re-generate the *preceding* words that “cornered” the LM into a position without a faithful continuation. In our example, we replace the token *in* by the token *with*; consequently, based on the context, the generated sentence can be completed with *my dog*.

Our method *Correction with Backtracking (CoBa)*, is a simple inference-time method that requires no additional model training and is compatible with most of the decoding methods. We evaluate CoBa on three established document summarization datasets and measure the faithfulness of generated summaries. We show that it is highly effective and efficient for detecting and mitigating hallucinations. CoBa is also orthogonal to many existing hallucination reduction techniques and can be used in conjunction to those.

\*Part of the work is conducted during an internship at Google.

### Context Document

I share my home with a loyal and affectionate companion - my dog. Living together has brought joy, companionship, and a unique bond into my life. Each day is marked by our shared adventures, whether it's going for long walks, playing fetch in the park, or simply enjoying quiet moments at home. Their unwavering presence brings comfort and a sense of connection, making every day brighter and more fulfilling.

Greedy Decoding: I live in **New York** . ✗  
Greedy+Backtrack: I live with **my dog**. ✓




Figure 1: Schematic illustration of CoBa (using only token probability as the detection metric with threshold 0.2). After the partial summary “I live”, the token “in” has a higher probability than “with”. However, “I live in” will pressure the model into hallucinating a place. We detect this because all the next tokens have a probability lower than our threshold 0.2. Backtracking enables the model to find an alternative continuation that avoids hallucination down the line.

## 2 Background and Related Work

We adopt the definition of hallucination for abstractive summarization from Maynez et al. (2020a): *The summary  $S$  for a context document  $C$  contains hallucinations if there exists a span in  $S$  which is not supported by  $C$ .*

Hallucinations exhibit task-specific characteristics in various Natural Language Generation (NLG) tasks. For instance, in Machine Translation, hallucination is often observed in the output when the input source undergoes specific perturbation (Lee

et al., 2018). In Question Answering (QA), one common manifestation is semantic drift, where the generated answers deviate from the topic of the question (Li et al., 2021). Additionally, in retrieval-based QA, the retrieval model may introduce additional sources of hallucination (Ji et al., 2023).

Various existing works seek to understand how hallucination happens, and have identified several factors. In various datasets, human generated ground truth summaries can contain additional information not present in the corresponding input texts (Dhingra et al., 2019; Wang et al., 2020). Training on such data may increase a model’s tendency to hallucinate. During generation, hallucination may occur when the model attends to irrelevant parts of the input context (Tian et al., 2019), or utilizes knowledge acquired during training that is not grounded in the context (Longpre et al., 2021). Additionally, the decoding method also impacts the faithfulness of generation. Past work has observed that sampling-based decoding can lead to increased hallucination (Dziri et al., 2021; Lee et al., 2022; Wan et al., 2023).

### 2.1 Methods for Reducing Hallucination

Depending on the task and problem setup, various methods have been developed to detect and mitigate hallucinations. Existing approaches can be broadly categorized into training time mitigation and generation time mitigation.

**Training Time Mitigation.** Noise in the pre-training corpus is shown to be a significant source of hallucination for language models (Zhou et al., 2023). Some past work has focused on applying simple mechanisms to filter training data, many of which are already used in training large language models (Touvron et al., 2023b; Penedo et al., 2023; Li et al., 2023b). Data curation is not only done in the pre-training stage but also can happen during supervised finetuning (SFT). Researches in this area focus on using high-quality, human curated, or domain-specific data (Elaraby et al., 2023) for SFT and have shown that this can lead to improved faithfulness (Zhou et al., 2023; Chen et al., 2023; Lee et al., 2023; Cao et al., 2023).

**Generation Time Mitigation.** Recent publications have also explored how to enhance the faithfulness of generation during inference time (Zhang et al., 2023). One line of work performs post-

editing by training specialized models (Cao et al., 2020; Chen et al., 2021; Dong et al., 2020) or by directly prompting the models (Varshney et al., 2023; Mündler et al., 2023). Others modify the decoding algorithm. Lee et al. (2022) proposes to gradually decrease the value of  $p$  in top- $p$  sampling (i.e. nucleus sampling), to reduce hallucinations introduced by randomness. Li et al. (2023a) modifies attention to encourage more factual generations. Shi et al. (2023) proposes *Context-Aware Decoding (CAD)* to suppress hallucinations arising from the model’s prior knowledge; they adjust the context-conditional token logits with the unconditional logits. Wan et al. (2023) proposes *Lookahead*: At each decoding step, it rolls out future summaries for the top  $k$  tokens with the highest probabilities, adjusts their probabilities with BS-Fact, and picks the token with the highest adjusted probability. They also show that the performance can be further improved by ranking multiple candidates with a composite faithfulness score, or by distilling student models with the generated summaries. In contrast to these methods, CoBa does not tamper with token probabilities. Instead, it detects hallucinated tokens and fixes them through backtracking and local edits (see Figure 1).

Most similar to our work is arguably King et al. (2022), a publication that we were not aware of until after the completion of this paper. While we do have distinct design choices and evaluations, we acknowledge that the two methods are rather similar and expect them to perform similarly under our setting.

### 3 Problem Setup

Let  $\mathcal{M}_\theta$  be an autoregressive summarization model with parameters  $\theta$ , and let  $\Sigma$  be its vocabulary. Given a context document  $\mathcal{C} = (c_1, \dots, c_m)$  as input,  $\mathcal{M}_\theta$  produces a summary  $\mathcal{S} = (s_1, \dots, s_n)$ :

$$\mathcal{M}_\theta(\mathcal{C}) = \mathcal{S}$$

where  $c_1, \dots, c_m, s_1, \dots, s_n \in \Sigma$ ;  $m$  and  $n$  are the lengths of the context and the summary respectively. In practice,  $\mathcal{M}_\theta$  can either be a specialized summarization model like PEGASUS (Zhang et al., 2020), or a general language model capable of zero-shot summarization like Flan-T5 (Chung et al., 2022). If  $\mathcal{M}_\theta$  requires prompting, we add a prompt like “summarize: ” to the context as input.




Figure 2: **Average token probability (top) and token-to-context distance (bottom) around the hallucination span.** Token offset 0 stands for the token where hallucination starts, negative offsets stand for the tokens before hallucination and positive ones are for the hallucinated tokens. On average, the token which starts the hallucination has the lowest probability and is the furthest away from the context tokens compared to surrounding ones.

Model  $\mathcal{M}_\theta$  generates the summary autoregressively. At each step, given a partially generated summary  $\mathcal{S}_{<t}$  up to token  $s_{t-1}$ , it outputs a distribution  $p_\theta(s_t|\mathcal{C}, \mathcal{S}_{<t})$  for the next token  $s_t$  over the vocabulary  $\Sigma$ . The probability of generating the summary  $\mathcal{S}$  is thus

$$p(\mathcal{S}) = \prod_{t=1}^{|\mathcal{S}|} p_\theta(s_t|\mathcal{C}, \mathcal{S}_{<t})$$

### 4 Reducing Hallucination at Inference

We present a detection-correction approach for reducing hallucination at decoding time. The main idea is illustrated in Figure 1: If a hallucination occurs, the problem typically originates already with its preceding tokens. The partially decoded summary can “corner” the model such that there is no faithful next token. For example, in Figure 1, the natural continuation for the partial summary “I

live in" is a name of a place. The source context however does not mention any places. We design strategies to detect such occurrences, and use backtracking (Tarjan, 1972) to find alternative phrases that prevent hallucinations down the line.

## 4.1 Hallucination Detection

We investigate different properties of hallucinated text and devise two strategies for detecting text that is not grounded in the context.

### 4.1.1 Uncertainty-based Detection

The intuition behind uncertainty-based detection is that hallucination is likely to occur if the model is unsure about what it should generate next conditioning on the input. The conditional probability of a token is one way of measuring uncertainty and prior work has shown that the token-wise probability of autoregressive language models is well-calibrated (Kadavath et al., 2022). Petryk et al. (2023) also use a similar technique for evaluating and ranking the correctness of image captions.

We validate that token probabilities are effective for identifying hallucinated tokens in summaries by computing probabilities on an annotated hallucination dataset from Maynez et al. (2020b). The dataset contains generated summaries from different summarization models, such as finetuned BERT (Devlin et al., 2018), Pointer-Generator Model (See et al., 2017) and several more, with human annotations for hallucination spans. Figure 2 presents the conditional token probabilities of Flan-T5 XL around the hallucination span. Offset 0 represents where the hallucination starts, the negative offsets represent preceding tokens and the positive offsets represent successive tokens. In the figure, we observe a significant drop in token confidence at the start of hallucination. The average probability is only 0.2 in contrast with 0.5-0.6 for non-hallucinated tokens. The distribution of the probabilities is noisy shown as wide standard deviation in the figure, because of annotation noise and because some generated summaries can contain unnatural segments.

Therefore, measuring conditional token probability is one way of detecting the beginning of hallucinations during the decoding process, when all possible next tokens have low probability, it suggests the absence of a suitable candidate, and potentially signals the onset of hallucination. Formally, at step  $t$ ,

we flag the token if the following condition holds:

$$p_{\theta}(s_t|\mathcal{C}, \mathcal{S}_{<t}) < \delta$$

where  $\mathcal{C}$  is the context document,  $\mathcal{S}_{<t}$  is the partially generated summary, and  $\delta$  is the token level conditional probability threshold for hallucination.

### 4.1.2 Similarity-based Detection

Another intuitive way of detecting hallucination is to find tokens in the generated summary that are not supported by the context, i.e., tokens that are not "close" to any part of the context document. One method of measuring closeness is by computing cosine distance in the embedding space of a language model. More concretely, given a proposed token, we compute the distance between its embedding and the embeddings of all tokens in the context and flag the token as a potential hallucination if the minimum distance is above a certain threshold. The detection criterion in this case is:

$$d(v, \mathcal{C}) = \min_{c_i \in \mathcal{C}} \cos\_dist(\text{Emb}(v), \text{Emb}(c_i)) > \varphi$$

where  $v$  is the proposed token,  $\mathcal{C}$  is the context document and  $\varphi$  is the distance threshold. Figure 2 presents the minimum token-to-context distance computed over the annotated dataset from Maynez et al. (2020b)'s with embeddings from Flan-t5 XL (the results are averaged over 5000 samples). The average token distance at the first word in a hallucination span is significantly higher than words at other positions, as expected.

## 4.2 Hallucination Mitigation

After detecting potential hallucination during decoding using the techniques described in subsection 4.1, we perform a local intervention to prevent the generation of hallucinated phrases. Specifically, we introduce a process similar to depth first search. We eliminate the last generated token  $s_t$  and try to propose an alternative token  $s'_t$  that does *not* satisfy the hallucination criteria. We keep track of the eliminations given a partial sequence  $\mathcal{S}_{<t}$  and context  $\mathcal{C}$  to avoid repetitive proposals. If  $s'_t$  can be found we add it to the generation and continue the forward decoding. We also continue if the partial sequence  $\mathcal{S}_{<t}$  only contains the start-of-sequence token [SOS]. Otherwise, we backtrack again, i.e. eliminate the current last token  $s_{t-1}$  and repeat the process (see Figure 1 for a pictorial description).



Admittedly, sometimes the model is unable to find a good solution, and this is signaled by backtracking too many times. We therefore introduce an upper bound  $L$ , for the number of decoding steps (both forward and backtracking) that can be performed. We pick  $L = 10T$  where  $T$  is the maximum generation length for our model  $\mathcal{M}_\theta$ . If an acceptable summary cannot be generated in  $L$  steps, we turn off the backtracking mechanism and adopt greedy decoding to generate the summary. We empirically observe that with reasonable threshold choices, less than 3% of the generations exceed the upper bound  $L$  when using moderate threshold values in general.

## 5 Experiments

### 5.1 Datasets and Models

We consider two models: Flan-T5 XL (Chung et al., 2022) and LLaMA (Touvron et al., 2023a). We use the pretrained models without any further finetuning on individual datasets.

We consider three datasets: Newsroom (Grusky et al., 2018), CNN/Dailymail (Nallapati et al., 2016), and XSUM (Narayan et al., 2018). We report numbers on the full test set of CNN/Dailymail and XSUM, and randomly sample a subset of size 5000 from the Newsroom test set. The XSUM dataset uses the first sentence of the original article as the ground truth summary, and the rest of the article as the context document. Consequently, core information is sometimes missing from the context. To improve the completeness of the context and enable more meaningful comparison with the ground truth, we adopt a similar approach as Wang et al. (2020) and prepend the ground truth summary back to the articles before performing summarization.

### 5.2 Baselines and Implementation Details

We examine four baseline decoding methods: greedy decoding, nucleus sampling, Lookahead (Wan et al., 2023) (see section 2), and CAD (Shi et al., 2023). Note that Lookahead takes a long time to roll out future summaries and compute BS-Fact for each of the rollouts (for instance, generating 5000 Newsroom samples takes 108 hours). One natural way of increasing the speeding of this method is to perform “lookahead” once every  $l$  tokens instead of after every token. Thus, we consider four choices of  $l$  for Lookahead:  $l = 1$  (the original version),  $l = 2$ ,  $l = 4$  and  $l = 8$ .

Additional implementation details can be found in subsection A.1 in the Appendix.

We consider two versions of CoBa: (1) CoBa that only uses the conditional word probabilities for detection, which we refer as CoBa in the tables; (2) CoBa that uses both the conditional word probability and the token-context distance, which we refer as CoBa-d. We use probability threshold  $\delta = 0.2$  and distance threshold  $\varphi = 0.5$  for Flan-T5, and  $\delta = 0.3$  and  $\varphi = 0.9$  for LLaMA.

We evaluate CoBa’s performance with greedy decoding and nucleus sampling. Since CoBa is complementary to most decoding methods, we can also use CoBa in conjunction with some of the baselines. We report results of using CoBa and CAD together. We do not evaluate using CoBa with Lookahead due to the high computational cost.

### 5.3 Metrics

To evaluate faithfulness, we compare the generated summaries with their *source documents*. We use the following metrics: **AlignScore** (Zha et al., 2023) and **FactCC** (Kryściński et al., 2019), both of which employ learned models to score faithfulness; **BS-Fact**, which measures the BERTScore (Zhang et al., 2019) precision of a generated summary with respect to its context document; **ROUGE-L** (Lin, 2004), which measures the longest common subsequence between the generation and reference. These metrics align relatively well with human judgement (Pagnoni et al., 2021b) and have reasonable runtime.

We also report standard summarization metrics, including **ROUGE-L**, **BERTScore F1** and **Bleurt** (Sellam et al., 2020), computed between the generated summaries and the datasets’ *ground truth summaries*. It should be noted that the models are used in a zero-shot manner. The quality of the generated summaries depends on the model’s capabilities, and they may have different styles compared to the ground truth. Therefore, this comparison may not always yield informative results.

### 5.4 Results

We report the faithfulness performance of Flan-T5 on the different datasets in Table 1, and the performance of LLaMA in Table 2. Note that all metrics are computed between the source document and the generated summary. We report the

Table 1: **Faithfulness of the summaries generated with various decoding methods using Flan-T5.** All the metrics are computed between the context document and the generated summary; higher is better.

	Method	AlignScore↑	FactCC↑	BS-Fact↑	Rouge-L↑
Newsroom	Greedy	0.765	0.604	0.919	0.131
	+ Lookahead (every 8 tok.)	0.768	0.607	0.920	0.133
	+ Lookahead (every 4 tok.)	0.774	0.607	0.922	0.136
	+ Lookahead (every 2 tok.)	0.811	0.662	0.931	0.153
	+ Lookahead (every tok.)	0.816	0.662	0.933	0.159
	+ CAD	0.746	0.490	0.916	0.145
	+ CoBa	0.821	0.674	0.923	0.138
	+ CoBa-d	0.865	0.709	0.926	0.145
	+ CoBa + CAD	0.773	0.515	0.919	0.149
	+ CoBa-d + CAD	0.820	0.560	0.922	0.161
	Nucleus	0.636	0.482	0.902	0.101
	+ CAD	0.694	0.430	0.907	0.117
	+ CoBa	0.800	0.645	0.920	0.128
	+ CoBa-d	0.857	0.692	0.923	0.139
	+ CoBa + CAD	0.767	0.505	0.917	0.139
	+ CoBa-d + CAD	0.817	0.552	0.921	0.154
XSUM	Greedy	0.723	0.485	0.919	0.096
	+ Lookahead (every 8 tok.)	0.727	0.486	0.919	0.096
	+ Lookahead (every 4 tok.)	0.733	0.487	0.920	0.097
	+ Lookahead (every 2 tok.)	0.756	0.514	0.925	0.101
	+ Lookahead (every tok.)	0.767	0.524	0.926	0.102
	+ CAD	0.694	0.383	0.919	0.094
	+ CoBa	0.752	0.504	0.920	0.096
	+ CoBa-d	0.791	0.523	0.921	0.104
	+ CoBa + CAD	0.707	0.398	0.919	0.094
	+ CoBa-d + CAD	0.735	0.414	0.923	0.103
	Nucleus	0.545	0.364	0.902	0.082
	+ CAD	0.621	0.317	0.911	0.088
	+ CoBa	0.730	0.489	0.917	0.093
	+ CoBa-d	0.772	0.499	0.920	0.101
	+ CoBa + CAD	0.695	0.373	0.918	0.093
	+ CoBa-d + CAD	0.728	0.392	0.922	0.102
CNN/DM	Greedy	0.840	0.506	0.922	0.146
	+ Lookahead (every 8 tok.)	0.843	0.511	0.923	0.147
	+ Lookahead (every 4 tok.)	0.848	0.514	0.925	0.149
	+ Lookahead (every 2 tok.)	0.866	0.546	0.930	0.157
	+ Lookahead (every tok.)	0.874	0.561	0.932	0.162
	+ CAD	0.828	0.301	0.917	0.173
	+ CoBa	0.869	0.554	0.924	0.149
	+ CoBa-d	0.884	0.570	0.925	0.151
	+ CoBa + CAD	0.836	0.312	0.918	0.174
	+ CoBa-d + CAD	0.849	0.330	0.919	0.178
	Nucleus	0.706	0.310	0.907	0.122
	+ CAD	0.777	0.232	0.911	0.157
	+ CoBa	0.857	0.521	0.922	0.142
	+ CoBa-d	0.872	0.533	0.923	0.145
	+ CoBa + CAD	0.828	0.291	0.916	0.169
	+ CoBa-d + CAD	0.841	0.313	0.918	0.174

metrics between the generated and ground truth summaries in Table 4 and Table 5 in the Appendix. For Flan-T5, both Greedy with CoBa and Lookahead at every token are competitive across datasets and metrics. Lookahead is slightly better according to BS-Fact and ROUGE-L, but is significantly

slower as seen in Figure 3. Greedy with CoBa is comparable to Lookahead every 4 tokens and is still much faster. For LLaMA, CoBa also attains performance gain. The improvement is smaller as LLaMA produces more faithful summaries than Flan-T5. It is important to note that the absolute

Table 2: **Faithfulness of the summaries generated with various decoding methods using LLaMA.** All the metrics are computed between the context document and the generated summary; higher is better.

	Method	AlignScore↑	FactCC↑	BS-Fact↑	Rouge-L↑
Newsroom	Greedy	0.701	0.321	0.897	0.161
	+ CAD	0.706	0.247	0.910	0.170
	+ CoBa	0.715	0.328	0.906	0.162
	+ CoBa-d	0.729	0.335	0.906	0.164
XSUM	Greedy	0.798	0.406	0.931	0.221
	+ CAD	0.783	0.335	0.931	0.237
	+ CoBa	0.800	0.410	0.932	0.221
	+ CoBa-d	0.805	0.418	0.933	0.223
CNN/DM	Greedy	0.750	0.316	0.900	0.152
	+ CAD	0.740	0.251	0.919	0.176
	+ CoBa	0.753	0.323	0.902	0.153
	+ CoBa-d	0.759	0.327	0.902	0.154

values of FactCC is smaller for LLaMA, because LLaMA produces much longer summaries than Flan-T5, while FactCC has negative correlation with the summary length. We report the distribution of generated summary length in Figure 6 in the Appendix, to show that the performance gain is not caused by producing shorter summaries.

In Figure 5, we present two qualitative examples comparing greedy decoding vs. CoBa and CoBa-d. In the first example, the greedy decoding produces the summary "The Boston Globe's review of "Looper" by John Sutter." with a name that does not appear in the source document. Backtracking successfully replaces it with the correct name. In the second example, although the extended name of the soccer club can include "United" based on real world knowledge, the document itself only refers the soccer club as "Scunthorpe". CoBa-d is able to detect this and remove "United".

## 5.5 Analysis

**Token Probability Threshold.** We examine the effects of using different values for the token confidence threshold, and present the results in Figure 4. We use the newsroom dataset and the Flan-T5 XL model. To better capture faithfulness, all the metrics are computed between source document and the generated summary. High value is better for all metrics. For AlignScore and BS-Fact, the improvement saturates at threshold 0.2-0.25, while FactCC continue to improve.

**Embedding Distance Threshold.** We perform ablation studies on the choice of embedding dis-

tance threshold. Intuitively, the smaller the distance threshold is, the more similar the generated summaries are to their original documents. Results are presented in Table 3. "N/A" represents not applying the embedding distance threshold. We use token probability threshold 0.2, the Newsroom dataset, and the Flan-T5 XL model for the ablation experiments. Decreasing the threshold improves the performance, and the improvement saturates around threshold 0.5.




Figure 3: **AlignScore vs. Generation Time.** Note that the the x-axis is in log scale. The curve for Lookahead represents doing lookahead every  $k$  tokens for  $k$  from 200 to 1. CoBa attains the highest AlignScore with more than 10x speedup.




Figure 4: **Ablation on the token confidence threshold for CoBa.** High is better for all metrics. Most metrics saturate around threshold 0.2-0.25.

**Document:** Dear God, what have they done to Joseph Gordon-Levitt's face? In Rian Johnson's time-travel action drama, "Looper," the star has been prosthetically altered with a fresh set of eyebrows, a snubbed nose, green contact lenses, and what appear to be new lips. He's supposed to look like the young Bruce Willis. What he resembles, mostly, is mid-period Devo. The imposture is hardly convincing, but "Looper" is fast enough, weird enough, and just about smart enough to make you forget about that. The movie wants to mess with your head, depositing us in a decrepit, overcrowded 2044 and then sprinkling on time-loop ... This is an article preview. The full story is available to BostonGlobe.com subscribers.

**Greedy Decoding:** The Boston Globe's review of "Looper" by **John Sutter**.

**CoBa:** The Boston Globe reviews "Looper" by **Rian Johnson**.

**Document:** Scunthorpe midfielder Neal Bishop has signed a one-year contract extension. The 35-year-old joined the Iron from Blackpool in 2013 and has made 119 league appearances for the League One side. He helped them to a third-placed finish this season, before they were beaten by Millwall in the play-off semi-finals. Bishop told the club website: "With the way the season finished, it's a sense of unfinished business and it was disappointing for all of us."

**Greedy Decoding:** Scunthorpe **United** midfielder Neal Bishop has signed a new one-year contract with the Iron.

**CoBa-d:** Scunthorpe have signed midfielder Neal Bishop on a one-year contract extension.

Figure 5: **Qualitative examples of greedy decoding vs. CoBa and CoBa-d.** The hallucinated content is marked in **red** and the corrected details are marked in **green**. CoBa and CoBa-d correctly remove the hallucinated content by triggering backtracking at corresponding positions and generate summaries with more and faithful details.

Table 3: **Ablation on the threshold on token embedding distance.** We use token confidence threshold  $\delta = 0.2$  while varying the distance threshold  $\varphi$  for all the experiments in this table.

Dist. Thresh	AlignScore $\uparrow$	FactCC $\uparrow$	BS-Fact $\uparrow$	Rouge-L $\uparrow$
N/A	0.821	0.674	0.923	0.138
0.9	0.825	0.677	0.924	0.139
0.7	0.859	0.699	0.925	0.143
0.5	0.865	0.709	0.926	0.145
0.3	0.867	0.718	0.920	0.146
0.1	0.867	0.720	0.920	0.146

## 6 Limitations and Future Work

In this study, we propose a method for reducing hallucinations in text summarization by backtracking. Our method consists of two steps: detection and backtracking. We employ two token level conditional probabilities and distance between generated tokens and context tokens to detect hallucinations. Both of these are effective ways of detecting hallucinated text, but there could be other complementary metrics that could improve detection. We defer the exploration of alternative metrics to future research endeavors.

While our primary focus in this paper is summarization models, our method can easily be extended to other applications where generating factual text

is paramount. For instance, in question-answering systems which first retrieve relevant documents and then generate an answer, we can define the retrieved documents to be the context and employ CoBa to produce factually correct answers.

## 7 Conclusion

Current decoding methods don't explicitly allow a model to re-generate some part of the generated text when there is no highly probable completion to the partial text. Such a scenario would lead to hallucinations because the model is uncertain about how to complete the sentence and will sample a low probability word. We show that there is a relatively simple solution to mitigate hallucination, which we refer to as Correction with Backtracking (CoBa). CoBa is an inference-time method that requires no additional models, is computationally efficient, and can be directly applied to diverse summarization models without retraining. CoBa detects hallucinations by using conditional probabilities of the generated tokens and measuring the distance between the generated text and the context. To correct the hallucinated text, it applies backtracking to before the hallucination and re-generates text to avoid ending up in positions with only low scoring token options. We empirically verify that CoBa is able to identify and rectify hallucinated tokens



during autoregressive decoding, and we show that CoBa produces more factual summaries for various datasets. Our future work includes exploring other detection strategies as well as extending CoBa to more diverse tasks.

## Acknowledgement

This research is supported by a gift from the Simons Foundation, grants from the National Science Foundation NSF (IIS-2107161, III1526012, IIS-1149882, OAC-2118310), Natural Sciences and Engineering Research Council of Canada (NSERC 567916) and IIS-1724282), the Cornell Center for Materials Research with funding from the NSF MRSEC program (DMR1719875), LinkedIn and NewYork-Presbyterian Hospital.

## References

- Meng Cao, Yue Dong, Jiapeng Wu, and Jackie Chi Kit Cheung. 2020. Factual error correction for abstractive summarization models. *arXiv preprint arXiv:2010.08712*.
- Yihan Cao, Yanbin Kang, and Lichao Sun. 2023. Instruction mining: High-quality instruction data selection for large language models. *arXiv preprint arXiv:2307.06290*.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, et al. 2023. Alpapasus: Training a better alpaca with fewer data. *arXiv preprint arXiv:2307.08701*.
- Sihao Chen, Fan Zhang, Kazuo Sone, and Dan Roth. 2021. Improving faithfulness in abstractive summarization with contrast candidate generation and selection. *arXiv preprint arXiv:2104.09061*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Bhuvan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William W Cohen. 2019. Handling divergent reference texts when evaluating table-to-text generation. *arXiv preprint arXiv:1906.01081*.
- Yue Dong, Shuohang Wang, Zhe Gan, Yu Cheng, Jackie Chi Kit Cheung, and Jingjing Liu. 2020. Multi-fact correction in abstractive text summarization. *arXiv preprint arXiv:2010.02443*.
- Nouha Dziri, Andrea Madotto, Osmar Zaiane, and Avishek Joey Bose. 2021. Neural path hunter: Reducing hallucination in dialogue systems via path grounding. *arXiv preprint arXiv:2104.08455*.
- Mohamed Elaraby, Mengyin Lu, Jacob Dunn, Xueying Zhang, Yu Wang, and Shizhu Liu. 2023. Halo: Estimation and reduction of hallucinations in open-source weak large language models. *arXiv preprint arXiv:2308.11764*.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies. *arXiv preprint arXiv:1804.11283*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Daniel King, Zejiang Shen, Nishant Subramani, Daniel S Weld, Iz Beltagy, and Doug Downey. 2022. Don’t say what you don’t know: Improving the consistency of abstractive summarization by constraining beam search. *arXiv preprint arXiv:2203.08436*.
- Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Evaluating the factual consistency of abstractive text summarization. *arXiv preprint arXiv:1910.12840*.
- Ariel N Lee, Cole J Hunter, and Nataniel Ruiz. 2023. Platypus: Quick, cheap, and powerful refinement of llms. *arXiv preprint arXiv:2308.07317*.
- Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fan-jiang, and David Sussillo. 2018. Hallucinations in neural machine translation.
- Nayeon Lee, Wei Ping, Peng Xu, Mostofa Patwary, Pascale N Fung, Mohammad Shoeybi, and Bryan Catanzaro. 2022. Factuality enhanced language models for open-ended text generation. *Advances in Neural Information Processing Systems*, 35:34586–34599.
- Chenliang Li, Bin Bi, Ming Yan, Wei Wang, and Songfang Huang. 2021. Addressing semantic drift in generative question answering with auxiliary extraction. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 942–947.

- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023a. Inference-time intervention: Eliciting truthful answers from a language model. *arXiv preprint arXiv:2306.03341*.
- Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023b. Textbooks are all you need ii: phi-1.5 technical report. *arXiv preprint arXiv:2309.05463*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. Entity-based knowledge conflicts in question answering. *arXiv preprint arXiv:2109.05052*.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020a. On faithfulness and factuality in abstractive summarization. *arXiv preprint arXiv:2005.00661*.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020b. On faithfulness and factuality in abstractive summarization. *arXiv preprint arXiv:2005.00661*.
- Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2023. Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation. *arXiv preprint arXiv:2305.15852*.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021a. Understanding factuality in abstractive summarization with frank: A benchmark for factuality metrics. *arXiv preprint arXiv:2104.13346*.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021b. Understanding factuality in abstractive summarization with frank: A benchmark for factuality metrics. *arXiv preprint arXiv:2104.13346*.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*.
- Suzanne Petryk, Spencer Whitehead, Joseph E Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. 2023. Simple token-level confidence improves caption correctness. *arXiv preprint arXiv:2305.07021*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. Bleurt: Learning robust metrics for text generation. *arXiv preprint arXiv:2004.04696*.
- Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Scott Wen-tau Yih. 2023. Trusting your evidence: Hallucinate less with context-aware decoding. *arXiv preprint arXiv:2305.14739*.
- Robert Tarjan. 1972. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2):146–160.
- Ran Tian, Shashi Narayan, Thibault Sellam, and Ankur P Parikh. 2019. Sticking to the facts: Confident decoding for faithful data-to-text generation. *arXiv preprint arXiv:1910.08684*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Neeraj Varshney, Wenlin Yao, Hongming Zhang, Jian-shu Chen, and Dong Yu. 2023. A stitch in time saves nine: Detecting and mitigating hallucinations of llms by validating low-confidence generation. *arXiv preprint arXiv:2307.03987*.
- David Wan, Mengwen Liu, Kathleen McKeown, Markus Dreyer, and Mohit Bansal. 2023. Faithfulness-aware decoding strategies for abstractive summarization. *arXiv preprint arXiv:2303.03278*.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. Asking and answering questions to evaluate the factual consistency of summaries. *arXiv preprint arXiv:2004.04228*.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. Alignscore: Evaluating factual consistency with a unified alignment function. *arXiv preprint arXiv:2305.16739*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.

Zheng Zhao, Shay B Cohen, and Bonnie Webber. 2020. Reducing quantity hallucinations in abstractive summarization. *arXiv preprint arXiv:2009.13312*.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206*.

## A Appendix

### A.1 Additional Implementation Details

For the baselines, we use the Hugging Face implementation<sup>1</sup> for greedy decoding and nucleus sampling, and the official code of lookahead<sup>2</sup>. We use our own implementation for CAD as we did not find existing publicly available implementation. We use  $\text{top-}p = 0.9$  for nucleus sampling. Lookahead performs rollout for the  $\text{top-}k$  tokens with the highest probabilities; we use  $k = 5$  following the original paper. CAD uses a scaling factor  $\alpha$  when adjusting the conditional probabilities with the unconditional probabilities; we use  $\alpha = 0.5$  following the original paper. During generation, we set the minimum generation length to be 2 and maximum generation length to be 200 for all decoding methods.

### A.2 Summarization Metrics between Ground Truth and Generated Summaries

We report the summarization metrics between the ground truth summaries and the generated summaries from Flan-T5 in Table 4 and LLaMA in Table 5. All decoding methods with both models demonstrate reasonable performance. It is important to note that the ground truth summaries for each dataset are collected by distinct criteria: CNN/Dailymail (Nallapati et al., 2016) uses the human-written story highlights in bullet points, XSUM takes the first sentence of a document (Narayan et al., 2018; Wang et al., 2020), and Newsroom uses the HTML metadata (Grusky et al., 2018). As the models are not further finetuned on the individual datasets, their summaries often exhibit different styles from the ground truth summaries. Consequently, the summarization metrics only provide limited insights into the quality of the generated summaries.

### A.3 Generated Summary Lengths

Figure 6 shows the lengths of generated summaries from Flan-T5 on the Newsroom dataset. In general, the length distribution is similar across different decoding methods.




Figure 6: **Number of tokens (left) and number of characters (right) in the generated summaries from Flan-T5 on the Newsroom dataset.** The lengths have similar distributions across generation methods.

<sup>1</sup><https://huggingface.co>

<sup>2</sup><https://github.com/amazon-science/faithful-summarization-generation>

Table 4: **Summarization metrics between the ground truth summaries from the dataset and the generated summaries using Flan-T5.** Higher is better.

	Method	ROUGE-L↑	BERTScore F1↑	Bleurt↑
Newsroom	Greedy	0.312	0.890	0.441
	+ Lookahead (every 8 tok.)	0.313	0.890	0.441
	+ Lookahead (every 4 tok.)	0.314	0.891	0.442
	+ Lookahead (every 2 tok.)	0.323	0.892	0.451
	+ Lookahead (every tok.)	0.322	0.892	0.451
	+ CAD	0.281	0.883	0.412
	+ CoBa	0.313	0.889	0.436
	+ CoBa-d	0.306	0.885	0.428
	+ CoBa + CAD	0.281	0.883	0.412
	+ CoBa-d + CAD	0.267	0.878	0.399
	Nucleus	0.267	0.883	0.406
	+ CAD	0.270	0.882	0.404
	+ CoBa	0.306	0.888	0.432
	+ CoBa-d	0.299	0.882	0.423
	+ CoBa + CAD	0.282	0.883	0.412
	+ CoBa-d + CAD	0.269	0.879	0.401
XSUM	Greedy	0.422	0.920	0.540
	+ Lookahead (every 8 tok.)	0.426	0.921	0.543
	+ Lookahead (every 4 tok.)	0.431	0.922	0.546
	+ Lookahead (every 2 tok.)	0.467	0.927	0.569
	+ Lookahead (every tok.)	0.483	0.929	0.578
	+ CAD	0.399	0.916	0.522
	+ CoBa	0.431	0.920	0.541
	+ CoBa-d	0.462	0.919	0.548
	+ CoBa + CAD	0.402	0.916	0.523
	+ CoBa-d + CAD	0.440	0.919	0.535
	Nucleus	0.297	0.902	0.459
	+ CAD	0.335	0.907	0.485
	+ CoBa	0.411	0.918	0.527
	+ CoBa-d	0.442	0.918	0.534
	+ CoBa + CAD	0.388	0.915	0.515
	+ CoBa-d + CAD	0.426	0.917	0.526
CNN/DM	Greedy	0.260	0.874	0.388
	+ Lookahead (every 8 tok.)	0.261	0.875	0.389
	+ Lookahead (every 4 tok.)	0.262	0.875	0.389
	+ Lookahead (every 2 tok.)	0.265	0.876	0.395
	+ Lookahead (every tok.)	0.265	0.876	0.396
	+ CAD	0.248	0.871	0.392
	+ CoBa	0.256	0.873	0.382
	+ CoBa-d	0.256	0.872	0.383
	+ CoBa + CAD	0.248	0.871	0.391
	+ CoBa-d + CAD	0.246	0.870	0.389
	Nucleus	0.235	0.870	0.370
	+ CAD	0.241	0.869	0.386
	+ CoBa	0.253	0.872	0.379
	+ CoBa-d	0.254	0.871	0.379
	+ CoBa + CAD	0.246	0.871	0.389
	+ CoBa-d + CAD	0.245	0.870	0.388



Table 5: **Summarization metrics between the ground truth summaries from the dataset and the generated summaries using LLaMA.** Higher is better.

	Method	ROUGE-L $\uparrow$	BERTScore F1 $\uparrow$	Bleurt $\uparrow$
Newsroom	Greedy	0.210	0.861	0.438
	+ CAD	0.207	0.872	0.429
	+ CoBa	0.212	0.870	0.439
	+ CoBa-d	0.214	0.868	0.436
XSUM	Greedy	0.376	0.915	0.564
	+ CAD	0.362	0.908	0.533
	+ CoBa	0.379	0.915	0.564
	+ CoBa-d	0.383	0.915	0.564
CNN/DM	Greedy	0.239	0.859	0.406
	+ CAD	0.236	0.872	0.401
	+ CoBa	0.240	0.861	0.407
	+ CoBa-d	0.240	0.860	0.405