

# Studying word order through iterative shuffling

Nikolay Malkin<sup>1,2,5</sup>

Sameera Lanka<sup>3</sup>

Pranav Goel<sup>4,5</sup>

Nebojsa Jojic<sup>5</sup>

<sup>1</sup>Mila, Université de Montréal

<sup>2</sup>Yale University

<sup>3</sup>Microsoft

<sup>4</sup>University of Maryland

<sup>5</sup>Microsoft Research

kolya.malkin@yale.edu, pgoel1@umd.edu, {sameera.lanka, jojic}@microsoft.com

## Abstract

As neural language models approach human performance on NLP benchmark tasks, their advances are widely seen as evidence of an increasingly complex understanding of syntax. This view rests upon a hypothesis that has not yet been empirically tested: that word order encodes meaning essential to performing these tasks. We refute this hypothesis in many cases: in the GLUE suite and in various genres of English text, the words in a sentence or phrase can rarely be permuted to form a phrase carrying substantially different information. Our surprising result relies on inference by iterative shuffling (IBIS), a novel, efficient procedure that finds the ordering of a bag of words having the highest likelihood under a fixed language model. IBIS can use any black-box model without additional training and is superior to existing word ordering algorithms. Coalescing our findings, we discuss how shuffling inference procedures such as IBIS can benefit language modeling and constrained generation.

## 1 Introduction

Is a model’s understanding of syntax a precondition for its understanding of natural language? Recent work on large language models (Devlin et al., 2019; Tenney et al., 2019; Rogers et al., 2021) has made this a popular hypothesis. Yet, models that consume only bag-of-words features but rival those that understand syntax have surprised researchers time and again (Iyyer et al., 2015; Joulin et al., 2017). New concerns have emerged that natural language understanding benchmarks may not be challenging enough to make sentence structure relevant (McCoy et al., 2019; Niven and Kao, 2019).

Syntax is an essential aspect of language (Chomsky, 1965). Sentence structure can be quite important: two sentences with very different meanings may use the same set of words (Fig. 1). But how much does syntax, as realized in word order, matter in typical English text? *Given the words that make*

The scared mouse chased the hungry cat.

The hungry cat chased the scared mouse.



Figure 1: Which word order is more likely?

*up a sentence, but not their order, is the order usually recoverable? If so, word order rarely encodes more information than is found in the bag of words.*

In the past, linguists could not have answered this question empirically. Manually ordering words into sentences is too laborious, and when there are multiple orders that satisfy grammatical constraints, one needs a way to choose among them.

With the power of large language models, we can reduce this question to a computational one and resolve both issues: *given the bag of words, find the word order that is most likely under a trained LM.* To make this search tractable, we develop inference by iterative shuffling (IBIS), a procedure inspired by techniques in combinatorial optimization, that is superior to existing approaches to this problem. Armed with IBIS, we answer the question above statistically and explore the implications.

First, we measure how often sentences and phrases are permutable in text of various genres.

Next, we analyze the effect of word order on the GLUE suite (Wang et al., 2018) and on the task of autoregressive language modeling. Randomly reordering input words drops the performance of models on nearly all tasks, but when we *infer* the order with the aid of a pretrained LM, this drop is small or absent. Thus, NLP pipelines can effectively consume bags of words as input, and order carries much less meaning than we might imagine.

We conclude with the implications of our results for language modeling. A computationally feasible search for word order clears a path for models that focus on content, rather than syntax, enabling a range of constrained generation applications.

NLL Ordered sentence	
<b>IBIS (10 search steps)</b>	
140.9	<u>the housed is 1967, period furniture the exterior beautiful in fine complemented by Restored</u> inside.
124.9	Restored is 1967, <u>period furniture the exterior beautiful in fine</u> complemented by the housed inside.
121.6	Restored is 1967, in fine <u>period furniture the beautiful exterior complemented by the</u> housed inside.
112.3	Restored <u>is 1967, in fine the</u> beautiful exterior complemented by the period furniture housed inside.
112.1	Restored in <u>fine is the 1967, beautiful exterior</u> complemented by the period furniture housed inside.
105.7	Restored in 1967, fine <u>beautiful exterior is the complemented by the</u> period furniture housed inside.
104.1	Restored in 1967, <u>fine exterior is complemented by the beautiful the</u> period furniture housed inside.
88.0	<i>Restored in 1967, the beautiful exterior is complemented by the fine period furniture housed inside.</i>
<b>Beam search (64) without future costs</b>	
116.9	is complemented by the beautiful exterior in the fine period furniture housed inside. Restored 1967,
<b>Beam search (64) with future costs</b>	
110.9	Restored exterior is complemented by the beautiful furniture housed in the fine 1967, period inside.

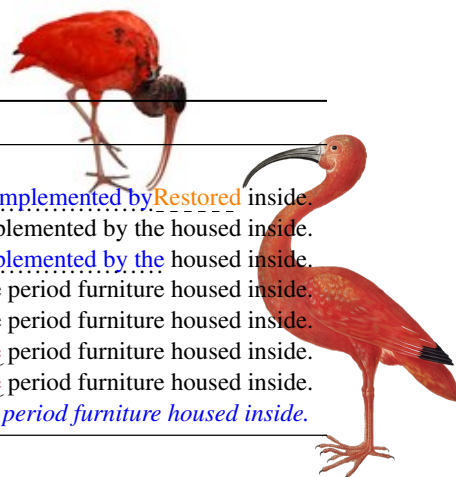


Figure 2: **Above:** IBIS iteratively infers the word order that has lowest negative log-likelihood (left column) under GPT-2. At each step, the sentence is cut into pieces, which are then rearranged. After several such  $k$ -opt moves, the *original order* is found. **Below:** Reconstructions of the same sentence using algorithms from prior work.

## 1.1 Related work

Research in cognitive science and psycholinguistics has raised the notion that syntax is a convention optimized for communicating bags of concepts over a linear channel. The emergence of syntactic phenomena is explained by information structure constraints (Jaeger, 2010; Levy and Jaeger, 2006; Hahn et al., 2020). Ours is the first large computational study to lend support to this view of syntax.

Circumstantial evidence for the redundancy of word order comes from work such as that of Niven and Kao (2019), which showed that language models’ predictions in certain tasks are largely explained by word-level triggers. Concurrently with this work, Sinha et al. (2021a,b); Pham et al. (2021); Gupta et al. (2021) probed and demonstrated, in various ways, the surprising insensitivity of infilling LMs’ performance on GLUE tasks to word order in training and evaluation data. These studies complement our discovery that nearly all of models’ accuracy on GLUE tasks can be explained by bags of words only (§5.2) to show that word order rarely carries information useful for classifying textual similarity, entailment, or sentiment.

In the domain of text generation, Khandelwal et al. (2018) found that the order of distant context words has little effect on prediction of the next word in a text. In §5.1 we confirm that the order of *recent* context strongly affects next-word prediction, but also show that this order can nearly always be inferred from the bag of words.

The problem of inferring word order from bags of words – *text linearization* – dates back to Elman (1990). This problem has been studied using both treelike and autoregressive models (de Gispert et al., 2014; Zhang and Clark, 2011; Liu and Zhang, 2015; Song et al., 2018). Horvat and Byrne (2014) reduce linearization under an  $n$ -gram model to generalized traveling salesman problems (TSP), but stop short of extending TSP algorithms to neural models, as we do in this work.

Algorithms based on best-first search were proposed by Liu et al. (2015) and Schmalz et al. (2016). The latter introduced a *beam search with future costs*, a key baseline in this paper. In the basic beam search algorithm for ordering a target bag of words, a LSTM model generates text from left to right, expanding a horizon of fixed size. The next-word distributions at each step are restricted to the words in the target bag that have not yet been used. The innovation of future costs is to modify the beam scoring function (which is usually the log-likelihood of the partially generated text) by adding the sum of log-likelihoods under a unigram model of the yet-unused tokens in the target bag.

## 2 Text linearization by iterative shuffling

IBIS is motivated by a need not only to generate text from left to right when inferring the most likely word order under a base LM, but to reason over an entire sentence and permute spans or words.

The bottom of Fig. 2 shows some of the fail-

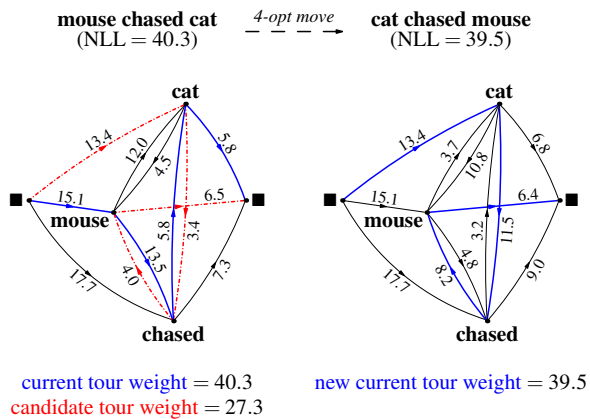


Figure 3: One IBIS search step: The edge weights in the left graph are negative log-likelihoods from GPT-2 Small conditioned on previous context in the current best sentence, “mouse chased cat” (thick blue path). A candidate 4-opt move that decreases the tour weight (dashed red path) is proposed and scored by the LM, yielding the graph on the right. Because the new NLL is lower than the old NLL, the move is accepted.

ure modes of beam search, with and without future costs. The reconstructions of a short sentence (in *italics*) by beam search using the GPT-2 Small model (Radford et al., 2019) are ungrammatical.<sup>1</sup> Beam search without future costs is unable to reason that the capitalized word is unlikely to occur in the middle of a text and should come first. Beam search with future costs suffers from the same inability to ‘plan ahead’: by the end of the sentence, the algorithm begins to fail as it is left with a set of words that cannot be arranged coherently.

However, some long spans that appear in the original sentence are generated, such as “is complemented by the” and “the fine period furniture housed inside”. The IBIS algorithm, which we will now describe, enables reasoning over the entire text to excise and recombine such coherent spans.

***k*-opt moves.** A *k*-opt move is the following operation: A sentence is ‘cut’ at *k* positions, creating *k* − 1 spans between the cuts. These *k* − 1 spans are then permuted (in one of (*k* − 1)! ways) to form a new sentence. Note that the cuts may come immediately before the first or after the last word.

Such operations were introduced by Lin and Kernighan (1973) in the context of the TSP on graphs. Recall that a *tour* of a weighted directed graph is a closed path that visits each vertex exactly once. A *k*-opt move on a tour is the operation of

<sup>1</sup>The beam size is set to 64 for illustration; with larger beam size, we observe similar failures for longer sentences.

removing *k* edges and inserting *k* new edges to create a new tour. Many heuristic algorithms for solving the TSP – finding the tour of lowest total weight – use *k*-opt moves as the core search step; a maximum of *k* = 5 is typical (Helsgaun, 2000).

There is a precise equivalence between text linearization and the TSP on graphs when the base language model is a bigram model. Suppose that the likelihood of a string has a factorization

$$p(w_0 w_1 \dots w_n) = \prod_{i=1}^n p(w_i | w_{i-1}), \quad (1)$$

where  $w_0$  and  $w_n$  are a fixed start/end token. We form a bidirected graph with vertices corresponding to the words of the sentence (and the start/end token) and set the weight of the edge from  $v$  to  $w$  to  $-\log p(w | v)$ . An ordering of the words is then equivalent to a tour of the graph, and its negative log-likelihood (NLL) is the weight of this tour. Finding the most likely order is equivalent to solving the TSP on this graph.

**Local search and the IBIS heuristic.** While *k*-opt search was developed with graph tours in mind, it can be applied to any scoring function, such as a language model, which produces a NLL of the next word depending on a long sequence of words preceding it,  $-\log p(w_{i+1} | w_0 w_1 \dots w_i)$ .

A naïve form of *k*-opt local search would find the order of a bag of words most likely under a base LM by beginning with a random candidate order, then repeatedly performing a random *k*-opt move, scoring the resulting order with the model, and accepting it as the new best candidate if it decreases the NLL. However, this approach is inefficient, as we will show below.

Instead, we propose a heuristic to improve the search. Let the current best order be  $w_0 w_1 \dots w_n$ . We form an auxiliary graph as above, but set the weight of the edge from  $w_i$  to  $w_j$  to

$$-\log p(w_j | w_0 w_1 \dots w_i), \quad (2)$$

that is, the NLL of word  $w_j$  at position  $i + 1$  given by the base LM. If the LM is a bigram model, then the (2) reduces to  $-\log p(w_j | w_i)$ , as before.

Now, we observe that the current order is a tour of this graph:  $w_0 \rightarrow w_1 \rightarrow \dots \rightarrow w_n$ . We rank all possible *k*-opt moves by how much they decrease the weight of this tour. Then, we create a batch of new candidate orders by performing *k*-opt moves sampled from near the top of this ranking

and we score this batch with the LM. The move that decreases the NLL most, if it exists, is accepted, yielding the new best candidate.

We emphasize that this is a heuristic, not an exact method. It is possible that a  $k$ -opt move decreases the weight of the tour in the auxiliary graph, but that when this move is performed and the sentence rescored by the LM, the NLL does not decrease. This is the case because the next-word probabilities given by the LM may depend on all preceding words. A  $k$ -opt move may change the context preceding a word  $w_i$ , which will modify the weights of the edges from  $w_i$  to other words in the graph. Nevertheless, the likelihood of a word depends mostly on recent context, especially on the preceding word, making IBIS an efficient heuristic.

**Practical considerations.** This heuristic for proposing  $k$ -opt moves is limited by computational constraints, mainly the difficulty of ranking all possible  $k$ -opt moves in memory. These difficulties arise in classical TSP solvers as well and are typically resolved by additional heuristics and sampling procedures. Our precise answers to these difficulties are described in Appendix A.

All our experiments are initialized with a random order of the target bag and use the heuristic described above to iteratively decrease the NLL. The search is terminated when there is no improvement for a specified number of steps (the ‘patience’ constant). For our experiments, we use a proposal batch size of 128 and a patience of 128 and limit the search to 3-, 4-, and 5-opt moves.

Our search for optimal  $k$ -opt moves uses core tensor operations and can run on a GPU – the first implementation of this kind, to the best of our knowledge. Runnable example code is provided in the associated repository: <https://github.com/malkin1729/ibis>. The reader can run the provided program and see a text of their choice iteratively shuffled into the most likely order (Fig. 5 in the Appendix).

### 3 Experiments: IBIS

We show that IBIS achieves a new state of the art in text linearization using a black-box LM. Following Wan et al. (2009), we evaluate on section 23 of the Penn Treebank (PTB) dataset of Wall Street Journal articles (Marcus et al., 1999).

**Setup and baselines.** The results here use GPT-2 Small as a base LM. There are three reasons for

Base LM	Inference	BLEU	NLL
LSTM	BS-512 +g	44.5	–
GPT-2 Small	BS-512	42.4	5.60
GPT-2 Small	BS-512 +g	45.4	5.48
GPT-2 Small	BS-1024	44.0	5.52
GPT-2 Small	BS-1024 +g	47.1	5.43
GPT-2 Small	IBIS	50.0	5.19
Distil-GPT-2	IBIS	47.2	–
GPT-2 Medium	IBIS	53.1	–
GPT-2 Large	IBIS	54.6	–
GPT-2 XL	IBIS*	53.6	–
GPT-2 Small	random $k$ -opt	44.1	5.36
–	original text	100.0	5.22

Table 1: BLEU scores of PTB sentences ordered by beam search and shuffling algorithms. NLL is negative log-likelihood per word under GPT-2 Small, averaged over test examples. First row is from Schmalz et al. (2016); +g denotes search with future costs. \*Batch size 64, patience 256 was used for GPT-2 XL.

this choice. First, it enables direct comparison between beam search and IBIS under the same base LM, which is difficult with the LSTM model of Schmalz et al. (2016) due to an incompatible code base. Second, it simultaneously allows us to measure the effect of the base model on results using the same inference procedure: IBIS is model-agnostic and can work with any base LM that produces next-token likelihoods. Third, GPT-2 is trained on generic English text and can handle arbitrary strings, making it a natural candidate for use in all parts of this paper. We chose the Small variant of the model for computational efficiency.

We used the published code of Schmalz et al. (2016), to tokenize the sentences, with minor processing for compatibility with GPT-2’s tokenizer. We then evaluated beam search with beam sizes 512 and 1024, with and without future costs. The future cost function used unigram frequencies estimated from the GPT-2 training data.<sup>2</sup> Any input word that is broken into multiple tokens by GPT-2’s subword tokenizer was always generated as a single unit. These baselines are thus directly comparable with the results of Schmalz et al. (2016).

**IBIS inference.** We ran the IBIS algorithm on this data, also using GPT-2 as the base LM, to infer an order of each sentence in the evaluation set.

<sup>2</sup><https://github.com/jhlau/acceptability-prediction-in-context>



To ensure that the bags of input and output words coincided, we did not allow  $k$ -opt moves that broke intact words between GPT-2 subword tokens.

**Results and discussion.** The BLEU scores of reconstructed sentences with respect to the original orders, and their NLLs per word, are shown in Table 1. IBIS outperforms beam search with size 512 and future costs – the strongest procedure in past work – by a large margin. Doubling the beam size, and the computation time, closes less than half of the gap in BLEU and in mean log-likelihood.

A less surprising, yet still meaningful, comparison is between base LMs: beam search with future costs with Schmalz et al. (2016)’s LSTM model and with GPT-2 Small (first and third rows). The former model is trained on a mix of target-domain (PTB) data, other datasets of news articles, and the Gigaword corpus, yet still resorts to using OOV tokens in place of infrequent words; GPT-2 is trained on a (biased) crawl of the Internet and processes rare words as sequences of subword tokens.

Finally, the mean log-likelihood per word under the base LM of sentences reconstructed with IBIS exceeds that of the original text. There are two ways to interpret this result. On one hand, it shows the strength of IBIS as an optimization algorithm: it may indeed be possible to permute the words in the original text, perhaps into an order more acceptable to human judgment, making word order more normative without changing the meaning. On the other hand, it shows that IBIS approaches the limit of what a text linearization algorithm that optimizes for GPT-2 Small likelihood can achieve as measured in BLEU score, which can be seen as a limitation of the base LM itself.

**Computation cost.** It is difficult to directly compare computation costs of beam search and IBIS due to the very different nature of these algorithms. A measure that corresponds well to the evaluation time is the number of calls to the base LM. During IBIS inference over the dataset of 2416 sentences, 67m strings were scored by GPT-2. Beam search with beam size 1024 and future costs would make approximately 60m calls to GPT-2 if all words were single tokens; handling of subwords increases this number to 76m. The computation time for the two algorithms was approximately equal. Thus IBIS is comparable to the BS-1024 +g baseline in computation, but performs significantly better.

Unlike beam search, which infers order incre-

Base LM	IBIS output
Distil-GPT-2	Ibises, all mud and crustaceans, usually feed as a group, have long downcurved bills, usually probing for food items.
GPT-2 Small	Ibises, all mud crustaceans, usually as a group, have long downcurved bills, usually probing for food items and feed.
GPT-2 Medium	Ibises as a group usually have long downcurved bills, usually probing for feed, mud, crustaceans, and all food items.
GPT-2 Large	Ibises usually feed as a group, usually have long, downcurved bills, all probing for food items, mud and crustaceans.
GPT-2 XL	Ibises usually feed as a group, and all have long, downcurved bills, probing for food items, usually mud crustaceans.
original	<i>Ibises all have long, downcurved bills, and usually feed as a group, probing mud for food items, usually crustaceans.</i>

Table 2: The words of a sentence from Wikipedia ordered by the IBIS algorithm with five base LMs.

mentally from left to right, IBIS works with the entire string at every step and can be stopped early to set a balance between time and output quality. As we note below, IBIS exceeds beam search in log-likelihood per word after far fewer search steps than were performed in our experiment.

**Dependence on base LM.** We tested IBIS on the PTB dataset using four other GPT-2 variants: the lighter Distil-GPT-2 and the larger GPT-2 Medium, Large, and XL<sup>3</sup>; the BLEU scores are shown in Table 1. More powerful base LMs improve the performance of IBIS, due to their greater ‘world knowledge’ or understanding of syntax, yet IBIS using the smallest model, Distil-GPT-2, still outperforms beam search with GPT-2 Small. Table 2 shows a sentence reordered using all five models.

**Importance of the heuristic.** We demonstrate the importance of the IBIS heuristic for proposing  $k$ -opt moves by performing the naïve  $k$ -opt local search described in §2 – randomly sampling  $k$ -opt moves, but keeping all other search parameters the same as for IBIS. On the PTB test data, the sentence reconstructions by this algorithm are significantly worse than those by IBIS: the search tends to exceed the patience (128 steps without improvement) at a higher negative log-likelihood.

<sup>3</sup>All pretrained models and weights acquired at: <https://github.com/huggingface/transformers>




Figure 4: The NLL/word on PTB sentences (defined as in Table 1) as a function of the number of  $k$ -opt batches, with and without the IBIS heuristic. IBIS tends to converge faster than an unguided  $k$ -opt search and stabilizes at lower NLL, yet both reach lower NLL than the computationally comparable beam search baseline (BS-1024+g).

Fig. 4 shows the mean negative log-likelihood per word as a function of the number of search steps, averaged over all sentences in PTB. IBIS reaches a lower NLL than the strongest beam search algorithm we evaluate after just 59 size-128 batches – equivalent to about  $\frac{1}{4}$  of the number of calls to GPT-2 made by beam search – and a lower NLL than the original text after 169 batches. Random  $k$ -opt search requires 6 times as many steps to reach the NLL of beam search.

Curiously, random  $k$ -opt search reaches a better NLL but a worse BLEU score than the beam search baseline, suggesting that beam search is good at correctly generating short spans of text (benefiting BLEU), but  $k$ -opt search is better at reasoning over the entire sentence (benefiting total likelihood).

#### 4 Experiments: (Im)permutability

Using the IBIS algorithm, we analyze the importance of word order in English text of different genres. We use three publicly available corpora covering different domains of textual expression:

**Yelp:** About 560k Yelp reviews, commonly used as a text classification benchmark.<sup>4</sup>

**Wiki:** 2m Wikipedia articles (Shaoul, 2010).<sup>5</sup>

**arXiv:** 1.7m scientific preprint abstracts (Clement et al., 2019)<sup>6</sup> filtered to remove  $\text{\TeX}$ .

IBIS can be applied to any bag of tokens, including punctuation marks, as shown in Fig. 2 and Table 2. In §3, we shuffled all tokens in a sentence to be consistent with the setup in prior work. However, to measure whether order is essential to conveying meaning, we face the problem of discerning whether punctuation marks are used to structure a compound thought or to separate distinct thoughts:

the pathological example is a stylistic choice to replace all periods with semicolons.

To simply investigate how long a phrase needs to be before it becomes possible to find significant rearrangements with higher likelihoods, we limit our analysis to text spans of two kinds: sentences that contain no punctuation and spans of text between two consecutive punctuation marks. We sample 1000 such sentences and spans from each of the three domains with lengths (in words) falling into each of several buckets (Table 4) and infer their most likely word orders using IBIS. For the spans between punctuation, 50 words of ordered context before the initial punctuation mark are provided for scoring of candidate word orders.

We analyze the reconstructions automatically using BLEU scores and via human evaluation.

**Results and discussion.** Table 4 shows the BLEU scores of the IBIS-inferred spans with respect to the original orders, as well as the ratio of perplexities under GPT-2 of the original and reconstructed texts. IBIS often finds orders that are more likely than the original ones in all three domains. The similarity of reconstructed and original spans at small lengths ( $< 30$  tokens) is remarkably high.<sup>7</sup>

We also see some differences between the domains. Especially at higher lengths, sentences and spans from Wikipedia and Yelp are difficult to permute into sentences with higher likelihood ( $\text{PR} > 1$ ). Reconstructed sentences from Yelp and arXiv retain fewer of the original 2-, 3-, and 4-grams (lower BLEU). Indeed, long Yelp sentences tend to ‘ramble’ using many frequent words, arXiv sentences are full of scientific terms that a non-expert can easily permute without losing grammat-

<sup>4</sup>[kaggle.com/ihamp31/yelp-review-dataset](https://kaggle.com/ihamp31/yelp-review-dataset)

<sup>5</sup>[www.psych.ualberta.ca/~westburylab/](http://www.psych.ualberta.ca/~westburylab/)

<sup>6</sup>[kaggle.com/Cornell-University/arxiv](https://kaggle.com/Cornell-University/arxiv)

<sup>7</sup>For comparison, machine translation systems rarely exceed 40.0 BLEU points; the score of IBIS on PTB (§3), where the average sentence length is 23 words, is 50.0.

*I was just there again in April and bragged to my friends about how great it was and we were all horribly disappointed.*  
 just bragged to all my friends about how great it was in April and we were there again and I was horribly disappointed.

*Mao Zedong’s philosophical essay furthered Marx and Lenin’s thesis and suggested that all existence is the result of contradiction.*  
 Marx and Lenin’s philosophical essay suggested and furthered Mao Zedong’s thesis that all existence is the result of contradiction.

*We introduce a longevity feature to the classical optimal dividend problem by adding a constraint on the time of ruin of the firm.*  
 We introduce the classical problem of the optimal ruin of a firm by adding longevity to the feature time constraint on a dividend.

Table 3: Examples of *original* and IBIS-reconstructed sentences from Yelp, Wikipedia, and arXiv (top to bottom). More examples are shown in Table 8.

(no punct.)	Yelp		Wiki		arXiv	
Length	BLEU	PR	BLEU	PR	BLEU	PR
5–9	81.5	0.94	77.4	0.91	75.0	0.91
10–19	69.5	0.91	67.4	0.90	51.6	0.80
20–29	51.3	0.97	51.9	0.96	37.0	0.79
30–39	35.7	1.04	38.8	1.01	27.2	0.79
40–49	24.7	1.01	31.4	1.07	23.1	0.82

(btw punct.)	Yelp		Wiki		arXiv	
Length	BLEU	PR	BLEU	PR	BLEU	PR
5–9	67.5	0.98	67.0	0.85	61.5	0.96
10–19	63.4	1.08	62.4	0.97	49.9	1.03
20–29	49.7	1.08	47.5	1.02	43.2	0.89

Table 4: Comparison of original and IBIS-inferred orders of punctuationless sentences (above) and spans between punctuation (below) of different lengths. We report BLEU score and ratio of perplexities (PR). A PR less than 1 indicates that IBIS reaches lower NLL per word than the original text.

icality, and Wikipedia sentences have a measured style more familiar to GPT-2 (see Table 3).

**Human evaluation.** Three human subjects were asked to rate the relationship of IBIS-inferred punctuationless sentences to the original texts. We sampled 50 sentences from each of the five length buckets from the Wiki dataset; annotators ranked each pair (original sentence, IBIS-inferred order) on a scale of 0 (the inferred order is unreadable or completely dissimilar to the original sentence) to 3 (the original and inferred orders are identical, achieved for 47 of the 250 sentences).

For punctuationless sentences with fewer than 20 words, more than half of pairs were given scores of 2 or 3 (similar or identical meaning). This number sharply drops with increasing length, but long punctuationless sentences are rare in normal text: most sentences without punctuation have fewer than 20 words. More details can be found in Appendix D.

$n$	perplexity			token acc %			$\pi$
	latent	top	rand	latent	top	rand	acc %
1		30.1			38.1		100.0
2	30.5	31.7	49.4	38.1	38.0	31.3	98.5
3	31.4	34.2	72.9	37.5	37.5	25.4	96.4
4	32.2	36.2	97.0	37.3	37.2	21.8	94.0
5	33.2	39.3	123.5	36.9	36.7	18.1	90.9
6	34.2	42.2	150.6	36.4	36.2	16.0	86.8
7	35.1	46.0	177.1	36.3	35.9	14.3	83.0

Table 5: Perplexity and accuracy of next-word predictions of GPT-2 Small, conditioned on the bag of the previous  $n$  words  $B$  and  $50 - n$  tokens of earlier context  $C$ . The last column is the frequency with which the true order of the  $n$  previous words is the most likely under the posterior  $p(\pi | C, B)$ .

## 5 Experiments: (Dis)order in NLP tasks

In this section we answer the question: How well could language models perform on standard NLP tasks if they were not given access to word order?

### 5.1 Word order and text generation

Left-to-right (autoregressive) text generation remains a principal direction of NLP research. How well could models such as GPT-2 generate text if, when prompted to generate the next word in a text, they did not know the order of the previous words? We measure the performance of GPT-2 in generating the next token in a text where the order of the previous  $n$  words is treated as a latent variable.

GPT-2 is a generative model of tokens, where preceding tokens (context) are used as predictor variables. We break this context into two parts: the distant ordered context  $C$ , followed by a bag of  $n$  tokens  $B = \{w_1, \dots, w_n\}$  whose order is not known (50 tokens in total). We evaluate the perplexity and word prediction accuracy of GPT-2 on a sample from OpenWebText, a reacquired version of the model’s training data, under three schemes for predicting the next word  $w_{n+1}$  given  $C$  and  $B$ :

**Latent.** There are  $n!$  possible orders of the bag

	entailment				paraphrase				
	CoLA	MNLI-m/mm	QNLI	RTE	MRPC	QQP	STS-B	SST-2	WNLI
	MCC	acc %	acc %	acc %	F1 / acc %	F1 / acc %	$\rho_P / \rho_S$	acc %	acc %
original	58.2	83.7 / 84.1	90.7	64.8	88.0 / 83.6	90.6 / 87.3	88.3 / 88.1	91.9	39.4
random	0.4	65.5 / 65.4	74.8	56.5	81.8 / 73.1	74.2 / 72.3	82.1 / 82.0	80.7	54.3
IBIS	39.2	79.4 / 79.8	86.4	64.8	86.1 / 81.0	88.9 / 84.4	87.8 / 87.5	86.0	40.3

Table 6: Standard metrics of finetuned BERT models (mean of 32 random seeds) on the GLUE benchmark tasks, evaluated on raw validation data, data with randomly ordered words, and data with word order inferred by IBIS.

of tokens  $B$ . We denote the order by  $\pi$ , a discrete latent variable taking values in permutations ( $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ ). By scoring *each* of these permutations following the context  $C$  under the base LM, we compute the posterior distribution over this latent variable conditioned on  $C$ :

$$p(\pi | C, B) \propto p_{\text{LM}}(w_{\pi(1)} \dots w_{\pi(n)} | C).$$

Under any order  $\pi$  of the past  $n$  tokens, the LM gives a distribution over the word  $w_{n+1}$ ,

$$p(w_{n+1} | \pi, C, B) = p_{\text{LM}}(w_{n+1} | C w_{\pi(1)} \dots w_{\pi(n)}).$$

In this setting, we predict  $w_{n+1}$  by integrating out the latent  $\pi$  (i.e., summing over all possible orders):

$$\begin{aligned} p(w_{n+1} | C, B) &= \sum_{\pi} p(\pi | C, B) p(w_{n+1} | \pi, C, B) \\ &\propto \sum_{\pi} p_{\text{LM}}(w_{\pi(1)} \dots w_{\pi(n)} w_{n+1} | C). \end{aligned} \quad (3)$$

**Top.** The same as **Latent**, but using only the top  $\pi$ , i.e.,  $p(w_{n+1} | \arg \max_{\pi} p(\pi | C, B), C, B)$ .

**Random.** In this case, we assume a uniform distribution over orders  $\pi$  of the bag  $B$  and predict

$$p(w_{n+1} | C, B) = \frac{1}{n!} \sum_{\pi} p(w_{n+1} | \pi, C, B). \quad (4)$$

This expression differs from (3) in that the likelihood under the base LM of the order of recent context is not taken into account: the order  $\pi$  of the bag  $B$  is assumed to be randomly sampled.

**Results and discussion.** The perplexity and token accuracy of GPT-2 Small under the **Latent**, **Top**, and **Random** schemes are shown in Table 5. Remarkably, even for 7 tokens of unordered context, integrating over a latent order reduces accuracy only about 2% from the model that has access to fully ordered context, and 83% of the time, the true order of the bag of 7 preceding tokens has the highest likelihood out of  $7! = 5040$  possible orders.

In light of the latter, it is unsurprising that the **Top** method has only lightly worse metrics than **Latent**.

However, the model rapidly degenerates when we randomly sample the order of the previous tokens. Indeed, the dependence of a word on a context word appearing  $m$  positions earlier sharply decreases with  $m$ . When the number of shuffled tokens  $n$  is large, the recent words, which are most predictive of the target word, are often moved far back in the context (see Appendix C for an example).

GPT-2 was trained with the objective of predicting a word given ordered context. We have shown that GPT-2 is able to infer the order of the context itself, then use it for prediction, while losing little in accuracy and perplexity. This is as much a result about language as it is a result about language models: the bag of tokens carries almost as much information as the ordered sequence.

## 5.2 Word order and GLUE

We evaluate the dependence of the GLUE benchmark (Wang et al., 2018) on word order in the input data. Because the inputs are often long, a full search over orders is infeasible, so we use IBIS.

Specifically, for each of the 9 tasks, we finetuned the BERT-Base model (Devlin et al., 2019), a standard baseline, on (ordered) training data using typical settings.<sup>8</sup> We then ran the IBIS algorithm with GPT-2 Small as the base LM to infer an order of the bag of words in each validation set sentence (in tasks with two input sentences per example, the sentences were ordered independently). The finetuned models were then evaluated on this IBIS-ordered validation data. For comparison, we evaluated the same models on validation data with words ordered randomly, as well as on the original orders.

<sup>8</sup><https://github.com/huggingface/transformers/blob/master/examples/pytorch/text-classification>



---

Fairy in techniques for **coronavirus vaccine** has been confirmed by research.

---

The mouse was hungry. It started feeding itself by taking UndergroundMISC’s engineered growth **cheese** assay. Now the mouse is still hungry.

---

The cat was hungry. Someone picked up the **mouse** and **chased** the cat away from it. Now the cat is still hungry.

---

The cat liked mice. His appetite for sweet treats was a little more intense. Now the cat hates mice.

---

The Dragon King has reached out to the court at the request of a judge with the Magic Kingdom. The Minister of Magic declined to comment.

---

Table 7: Constrained generation using IBIS variants (Appendix B): sentences were forced to begin or end with the underlined spans and to contain the **bold** words in any order; all other words were generated by the model.

**Results and discussion.** The standard evaluation metrics for these models are shown in Table 6. At least 95% of the prediction accuracy on tasks related to textual entailment (MNLI, QNLI, RTE), 97% on tasks evaluating similarity detection (MRPC, QQP)<sup>9</sup>, and 94% on the sentiment classification task (SST-2) is explained by bags of words alone. That is, *such high scores can be achieved by a model that consumes only bags of words as input*. Our model is a composition of a combinatorial search (IBIS) with a feedforward model (BERT), but these results place a lower bound on what models that are not given word order can achieve.

The last two tasks, CoLA and WNLI, do not follow this pattern. The anomaly of WNLI, which tests resolution of ambiguous anaphora, seems to be due to the tiny size of the data (71 validation examples); the baseline models, on average, perform worse than random guessing. On the other hand, CoLA tests grammaticality judgments, which clearly depend on word order; many examples have no grammatical order in the first place. The large drop in Matthews correlation is unsurprising.

The finetuned models perform substantially worse on evaluation data with randomly ordered words on all tasks (except WNLI), though still much better than chance (except on CoLA). We conclude that the trained models need word order to perform, but that the word order itself carries little information, as we can infer word orders that result in near-baseline evaluation scores.

## 6 Conclusion and future work

We have shown that word order in an English sentence encodes surprisingly little information in addition to that contained in the bag of words. NLP models such as BERT and GPT-2 depend on order when creating representations of text, because they were trained on ordered words, but at the same time do not strictly need it, since their understanding of syntax – and the compressed world knowledge that

<sup>9</sup>The metrics for STS-B are correlations, not accuracy.

they hold – are sufficient to infer word order.

It would be interesting to use techniques such as IBIS to study investigate humans’ capacities for syntax, both productive and receptive. Are sentences with unlikely word order – as measured by a language model – more likely to lead to confusion (as in the first and last rows of Table 8)?

A bolder conjecture states that many aspects of English syntax can be explained by optimality for language modeling. If, for a corpus of unordered sentences, we *jointly* infer a most likely word order for each example and a language model that fits these orders, do the inferred orders recover true English syntax, or at least a syntax satisfying known cross-lingual universals? If so, we would be led to vastly generalize the main claims of Levy and Jaeger (2006) and Hahn et al. (2020). An iterative ordering algorithm like IBIS is an essential step towards answering such questions.<sup>10</sup>

Our work can guide and motivate research into combining long-range dependencies in the evolution of content – vocabulary constraints such as sentiment, global story arcs, rules of rhyme and meter, etc. – with models like GPT-2 that are capable of generating and scoring text. As we discuss in Appendix B, variants of IBIS can be used for a wide variety of such constrained generation tasks by making some of the words in the bag latent and sampling them in concert with  $k$ -opt moves: generating text from keywords, constraining text to a fixed length, composing poetry, and others where beam search is inefficient (Table 7).

Thus, IBIS is an attractive, flexible alternative to beam search in generative language models. It may find applications well beyond word ordering.

<sup>10</sup>One can train a LM on initially unordered text, in an EM-like procedure that iteratively reinfers optimal word orders using IBIS and performs gradient steps. We considered this question with toy data and small models, but training a full GPT-2 in this way was far beyond our scope in computation costs and time. However, any empirical results in this direction, *especially* with small models, could have deep implications in linguistics.

## Acknowledgments

The authors thank the anonymous reviewers for their comments and suggestions.

## Ethics statement

We use this section not only to promote discussion of possible societal impacts, but also to help researchers keep certain things in mind when they look to use our method and results.

**Annotation Process.** All three human annotators (§4) have English as their native or first language and are at least college-graduated. They were all compensated at the rate of US\$15 per hour. They were made aware that the first of the two sentences they are shown is from English Wikipedia while the second sentence is a reordering that need not be grammatical. We share the full set of instructions given to help them do the rating task in Fig. 8.

**Use of Large Language Models.** The usage of large language models has significant environmental and financial impacts. However, the majority of the cost is borne by training a new large language model, rather than using an already trained *fixed* language model as we do in our algorithm. We posit that using more computationally feasible inference methods using existing large language models and resources, instead of training even bigger models or methods that require months of compute cycle, makes our work more usable and accessible.

Large pretrained language models are also known to carry significant social biases, and this might affect the optimal ordering of a bag of words that our approach may find (since our search space stems from the probabilities learned by large language models). In fact, it will be interesting to conduct a study specifically focused on language model preferences for word order in cases where the subject and object of a text operate in an imbalanced power hierarchy: we expect the training data of language models to have an impact on the recovered word order. On a broader note, our findings of how sentence structure may often be redundant in English text could be a premise for further work in sociolinguistics about variation in norms of word order.

**Language.** It is important to keep in mind that our experiments pertain to the English language and that our findings and implications should not

be transferred to other languages without further experimentation. For example, we may expect measures of permutability to differ significantly in synthetic languages that mark grammatical roles by suffixation and have a freer word order: a Russian or Warlpiri sentence is more likely to have a grammatical reordering than an English or Mandarin one, but this reordering may have nearly the same meaning as the original sentence, with the context of surrounding sentences playing a large role in conditioning topicalization. In other highly agglutinative languages, entire complex sentences can be expressed in a single (orthographic) word, and the very notion of ‘word’ as separate from ‘morpheme’ is difficult to define – a challenge for NLP models.

## References

- Noam Chomsky. 1965. *Aspects of the Theory of Syntax*. MIT Press.
- Colin B. Clement, Matthew Bierbaum, Kevin P. O’Keeffe, and Alexander A. Alemi. 2019. [On the use of arxiv as a dataset](#).
- Adrià de Gispert, Marcus Tomalin, and Bill Byrne. 2014. [Word ordering with phrase-based grammars](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 259–268, Gothenburg, Sweden. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jeffrey L. Elman. 1990. Finding structure in time. *Cognitive Science*, 14(2):179–211.
- Ashim Gupta, Giorgi Kvernadze, and Vivek Sriku-mar. 2021. [Bert & family eat word salad: Experiments with text understanding](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35:12946–12954.
- Michael Hahn, Dan Jurafsky, and Richard Futrell. 2020. Universals of word order reflect optimization of grammars for efficient communication. *Proceedings of the National Academy of Sciences*, 117(5):2347–2353.
- Keld Helsgaun. 2000. An effective implementation of the Lin–Kernighan traveling salesman heuristic. *European Journal of Operational Research*, 126(1):106–130.

- Matic Horvat and William Byrne. 2014. [A graph-based approach to string regeneration](#). In *Proceedings of the Student Research Workshop at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 85–95, Gothenburg, Sweden. Association for Computational Linguistics.
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. [Deep unordered composition rivals syntactic methods for text classification](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1681–1691, Beijing, China. Association for Computational Linguistics.
- T. Florian Jaeger. 2010. Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, 61(1):23–62.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Urvashi Khandelwal, He He, Peng Qi, and Dan Jurafsky. 2018. [Sharp nearby, fuzzy far away: How neural language models use context](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 284–294, Melbourne, Australia. Association for Computational Linguistics.
- Roger Levy and T. Florian Jaeger. 2006. Speakers optimize information density through syntactic reduction. In *Advances in Neural Information Processing Systems*, volume 19, pages 849–856.
- Shen Lin and Brian Kernighan. 1973. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21(2):498–516.
- Jiangming Liu and Yue Zhang. 2015. [An empirical comparison between n-gram and syntactic language models for word ordering](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 369–378, Lisbon, Portugal. Association for Computational Linguistics.
- Yijia Liu, Yue Zhang, Wanxiang Che, and Bing Qin. 2015. [Transition-based syntactic linearization](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 113–122, Denver, Colorado. Association for Computational Linguistics.
- Mitchell P Marcus, Beatrice Santorini, Mary Ann Marcinkiewicz, and Ann Taylor. 1999. *Treebank-3. Linguistic Data Consortium, Philadelphia*, 14.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Thang Pham, Trung Bui, Long Mai, and Anh Nguyen. 2021. [Out of order: How important is the sequential order of words in a sentence in natural language understanding tasks?](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1145–1160, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, R. Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. A Primer in BERTology: What We Know About How BERT Works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Allen Schmalz, Alexander M. Rush, and Stuart Shieber. 2016. [Word ordering without syntax](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2319–2324, Austin, Texas. Association for Computational Linguistics.
- Cyrus Shaoul. 2010. The Westbury Lab Wikipedia corpus. *Edmonton, AB: University of Alberta*, 131.
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021a. Masked language modeling and the distributional hypothesis: Order word matters pre-training for little. [arXiv:2104.06644](#).
- Koustuv Sinha, Prasanna Parthasarathi, Joelle Pineau, and Adina Williams. 2021b. [UnNatural Language Inference](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7329–7346, Online. Association for Computational Linguistics.
- Linfeng Song, Yue Zhang, and Daniel Gildea. 2018. [Neural transition-based syntactic linearization](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 431–440, Tilburg University, The Netherlands. Association for Computational Linguistics.

Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.

Stephen Wan, Mark Dras, Robert Dale, and Cécile Paris. 2009. [Improving grammaticality in statistical sentence generation: Introducing a dependency spanning tree algorithm with an argument satisfaction model](#). In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 852–860, Athens, Greece. Association for Computational Linguistics.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Yue Zhang and Stephen Clark. 2011. [Syntax-based grammaticality improvement using CCG and guided search](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1147–1157, Edinburgh, Scotland, UK. Association for Computational Linguistics.

## A More on IBIS

### A.1 Search parameters and code

We describe the search strategy of IBIS. As noted in the main text, a complete enumeration of  $k$ -opt moves to rank in the batch proposal step is not feasible. Thus we do the following:

- (1) At each step, we randomly sample  $k \in \{3, 4, 5\}$  and a permutation of  $(k - 1)!$  spans resulting from cutting the candidate sentence at  $k$  points. We search only for  $k$ -opt moves that permute the spans according to this permutation.
- (2) For long sentences it is impractical or infeasible, due to memory constraints, to compute the improvement in tour weight under *every*  $k$ -opt move – for a sentence of length  $N$ , the number of such moves is  $O(N^k)$ . Thus we sample a smaller set of candidate cut positions and score only  $k$ -opt moves that cut the sentence at positions in this set, alternating two strategies: (a) sampling 20 ( $k = 5$ ) or 40 ( $k = 4$ ) random candidate cut positions and (b) taking between 7 and 14 consecutive cut candidates at a random position in the text.

```
(base) [redacted] $ python shuffler.py
Loaded model gpt2
Enter a sentence (>5 words)...
I am a language model. Do I understand syntax?
Original order NLL = 55.122459411621094
0 init 85.22409 model language syntax am a. Do I understand I?
1 GS 4 72.14638 . Do I understand a model language syntax am I?
3 GS 5 67.34329 . Do I understand syntax am I a language model?
11 GS 3 66.11749 Do I understand syntax. am I a language model?
19 GS 3 64.38109 am I a language model? Do I understand syntax.
27 GS 4 61.405075 am I a language model. Do I understand syntax?
53 GS 3 55.12246 I am a language model. Do I understand syntax?
```

Figure 5: IBIS example code usage.




Figure 6: Number of IBIS search steps and number of accepted  $k$ -opt moves plotted with sentence length (in words) in the PTB dataset.

- (3) We rank all  $k$ -opt moves, with the given  $k$  and permutation of spans, that cut at the candidate positions by how much they improve the current tour of the auxiliary graph with its current weights. A random  $b$  of the top  $B$  moves are proposed as the candidate batch. (We chose  $B = 512$  and took  $b$ , the batch size, to be 128.)

Our experiments were run on a mixture of Nvidia Tesla K80 and P40 GPUs. The latter are able to run GPT-2 Large with batch size 128 on texts the length of the longest sentence in the PTB dataset.

Figure 6 shows the distribution of search steps and accepted  $k$ -opt moves by sentence length.

### A.2 Visualizations

The associated repository includes three videos showing the evolution of texts with a widely known order, and weight matrices of the auxiliary graphs, as IBIS shuffles them into their optimal orders.

### A.3 Permutation examples

Table 8 shows more examples of original sentences and the word orders restored by IBIS, meant to illustrate its various success and failure modes. Of note, the first example is commonly misinterpreted by humans as having the opposite meaning. It would be interesting to study whether difficulty in communication between humans arises when spoken sentences can be permuted into text that is more likely and has very different meaning.



<i>Let's make a bet: winner owes loser 50 dollars.</i>	Let's make a bet: loser owes winner 50 dollars.
<i>The mouse chases the cat.</i>	The cat chases the mouse.
<i>Thoughts without content are empty, intuitions without concepts are blind.</i>	Thoughts without content are empty, intuitions without concepts are blind.
<i>Experience without theory is blind, but theory without experience is mere intellectual play.</i>	Experience without theory is blind, but experience without theory is mere intellectual play.
<i>Heat 12 oz. light beer, 1/2 tsp. Dijon mustard on low; whisk in 4 c. shredded sharp cheddar cheese until melted and smooth.</i>	Heat 1/2 tsp. Dijon mustard, melted in 4 oz. beer; whisk 12 c. shredded cheddar cheese on low light until smooth and sharp.
<i>To be, or not to be, that is the question: whether 'tis nobler in the mind to suffer the slings and arrows of outrageous fortune, or to take arms against a sea of troubles, and by opposing end them: to die, to sleep; no more; and by a sleep, to say we end the heart-ache, and the thousand natural shocks that flesh is heir to?</i>	To say we suffer is to take a thousand shocks, whether in the flesh, or by the nobler heart, or by the more outrageous sea-ache and slings of arrows, and to be against them: and to end troubles, to end arms, to be of the opposing mind; and to die: no question, that 'tis not a sleep; that sleep is the natural heir to fortune?
<i>Remarkably, even for entire paragraphs, the heuristic search is able to find mostly grammatical and somewhat coherent orderings.</i>	Remarkably, heuristic search is mostly grammatical, somewhat coherent and even able to find the orderings for entire paragraphs.
<i>It certainly was cold, he concluded, as he rubbed his numbed nose and cheek-bones with his mittened hand.</i>	It was certainly cold, he concluded, as he numbed his hand and rubbed his nose with mitten-ed cheekbones.
<i>This much-needed paper fills a gap in the literature.</i>	This paper fills a much-needed gap in the literature.

Table 8: Miscellaneous *original sentences* and the rearrangements of their bags of words inferred by IBIS under GPT-2.

## B IBIS beyond linearization

In this section, we explore a few additional applications of IBIS and its variants. The generated examples presented here were chosen out of multiple runs for each prompt and not thoroughly evaluated by automatic metrics, but are rather intended to suggest possible uses and advantages of order-free generation.

### B.1 Latent bags of words and lexical constraints

Suppose that we aim to generate a text with lexical constraints; e.g., the number of instances or particular words or the total length of the text may be fixed. Under an autoregressive LM, sampling from the set of sequences where the constraints are satisfied is in general intractable; for instance, it

is even intractable to sample from the distribution over sequences of length 10 ending with a period.

The IBIS algorithm can be modified to search for the most likely sequences of tokens satisfying such constraints. For example, suppose we are generating a text with six tokens that is required to contain the words ‘cat’ and ‘mouse’. We initialize a search with a sentence with ‘cat’, ‘mouse’, and four random words. To the IBIS search step of permuting spans of text, we add a step of replacing any word – besides ‘cat’ and ‘mouse’ – with any other word in the vocabulary. A batch proposal heuristic is possible here as well: we sample candidate replacement words at a position from the (perhaps softened) distribution over words at this position under the base LM given the current context.

Table 9 shows examples of sentences inferred

by such a search, constrained to begin with certain words and to contain certain other words.

## B.2 Reverse generation.

IBIS is readily modified to fix last few words of the generated text by simply restricting the set of candidate cut positions for  $k$ -opt moves. Thus we can generate text constrained to end with given words. IBIS with a word replacement step enables a faster and more robust *reverse* generation using only a forward LM (Table 10).

Notice that this search is able to find strings relevant to *future* context: when the sentence is forced to end with a span about cats, replacing a word in the middle of the sentence with ‘cat’ increases the likelihood of the entire text. At some point in the search, ‘cat’ gets sampled as a replacement, and the search enters a NLL sink: the word ‘cat’ is now likely to remain.

## B.3 Rhyming constraints

To give a taste of what further applications are possible, we use a modified IBIS to generate short verses. Rhyme and meter are lexical constraints that can be incorporated into word replacement search steps: for example, the words sampled for replacement at certain positions may be required to lie in the set of words that rhyme with an already generated line.

Let us rewrite Shelley’s famous lines:

Rise like Lions after slumber  
In unvanquishable number—  
Shake your chains to earth like dew  
Which in sleep had fallen on you—  
Ye are many—they are few.

with the help of Distil-GPT-2, forced to keep the two underlined lines and generate two new rhyming lines of appropriate length:

If you have a ton of lumber  
In unvanquishable number—  
Then enjoy your lumber stew—  
Ye are many—they are few.

Similarly, the following haiku verses, the result of a human-machine collaboration between the authors and GPT-2, were constrained to use the bold words and satisfy metric constraints.

**Fuji.** Simple **frog**,  
humble, short feet, do you know  
the distance to **home**?

See the early **moon**  
Night watcher’s little **lantern**  
But I see nothing

## C On random and latent next-word prediction

The goal of this short section is to illustrate how randomly sampling an order of recent context degrades the performance of GPT-2 in predicting the next token, while inferring it as a latent order does not. Consider the input: “Our Father, who art in Heaven, hallowed be thy” (l indicates subword division). GPT-2 recognizes this standard text and would predict the next word as ‘name’ if given ordered context. Integrating over a latent order of the last  $n = 2, \dots, 7$  tokens, or taking a random order of the last 2 or 3 tokens, the most likely next token is still ‘name’. However, for  $n = 4$ , the most likely next token under a random order is ‘lowed’: the distribution over next tokens is quite flat, and the most significant pattern is that ‘hall’ is the last token in  $\frac{1}{4}$  of orders – strong evidence for ‘lowed’ to follow.

## D Human evaluation on punctuationless sentences

All pairs of punctuationless sentences (original and IBIS-inferred orderings) were independently rated by 3 annotators, where each annotator followed the instruction set laid out in Fig. 8.

The distribution of users’ ratings for punctuationless sentences in each of the length buckets is shown in Fig. 7. For sentences in the most common length bucket (10–19), almost half of sentences are either identical to or have the same or similar meaning as the original text.

12 11 10 9 8 7 6 5 4 3 2 1  
 -4 -3 -2 -1 0 1 2 3 4 5 6 7

Fairy in techniques for **coronavirus vaccine** has been confirmed by research.  
 In the aftermath of the storm, **humans** and **hippos, penguins**, dolphins, and the dinosaurs were forced into the **streets** by the hive mind of the Internet.  
 See the full video for a long list of **physicists** and **Mars'** history as well as insights to the origins of **bubbles**.  
 Drones, **drugs, climate change**: the search for answers to nothing.

Table 9: Examples of constrained generation using IBIS endowed with a word replacement step. Underlined words are a fixed initial context, while the bold words are required to appear anywhere in the text, in some order. The base LM was Distil-GPT-2; a slight relaxation of greedy ascent is employed.

The cat was hungry. Someone picked up the **mouse** and **chased** the cat away from it. Now the cat is still hungry.  
The mouse was hungry. It started feeding itself by taking UndergroundMISC's engineered growth **cheese** assay. Now the mouse is still hungry.  
The cat liked mice. His appetite for sweet treats was a little more intense. Now the cat hates mice.  
The world just had a cat and a dog fight. Now the cat isn't hungry anymore.  
The cat and the mouse were both hungry. Although they oversee the animal kingdom, their predators eat more. Now only one of them is hungry.  
 I thought I could find my future and fix it, but everyone was running in panic.  
 The announcement of the student strike is expected to be welcomed by many of us. Students across Canada are rejoicing.  
 The Dragon King has reached out to the court at the request of a judge with the Magic Kingdom. The Minister of Magic declined to comment.

Table 10: Examples of reverse generation with an IBIS variant. Just as in Table 9, the underlined tokens are fixed and the bold tokens are required to appear.




Figure 7: Annotators' ratings of similarity between original and reordered sentences (§D).

Rate or label the given pair of sentences for "the similarity in their meanings", where the label can be 0/1/2 as follows:

**LABEL 0:** The two sentences are essentially completely different with not even surface similarities in what they mean or how they are structured, etc. This includes cases where the second sentence by itself seems non-sensical or meaningless because it is completely ungrammatical or not how English sentences look like (according to your understanding) by starting the sentence from the middle, etc. Examples of this score:

Sentence 1: Tyrone Winsloe takes an interest in Elena, wanting her to wear skimpy clothing as well as watch, and participate in, his 'hunts'.

Sentence 2: hunts, Winsloe takes an interest in Elena, wanting to participate in and watch her, as well as wear his skimpy clothing 'Tyrone'.

Label: [0]

Sentence 1: When he would pass out, the staff would throw water on him.

Sentence 2: When staff would pass him on, he would throw out the water.

Label: [0]

Sentence 1: Always for Canale 5, Boldi has interpreted his first fiction tv, giving life to the character of "big dad", Lorenzo Fumagalli in the television series "Un ciclone in famiglia", directed by Carlo Vanzina.

Sentence 2: for tv series directed by Lorenzo F. Boldagalli, Carlo Van toi cicloneum, first television character in "Unbigale", has interpreted the life of his dad in "Always 5", giving the Canzina famiglia fiction

Label: [0]

**LABEL 1:** There are surface similarities between the two sentences or they are same except for some entities being switched around but in a way that the meaning changes. The meaning of the two sentences is different overall, though they say the same things in part. This can be loosely translated as "some similarities but different in their overall meaning". Examples:

Sentence 1: "In 1895 they joined the South East Lancashire League, and a new professional and coach was employed, a John Redfern from Linthwaite."

Sentence 2: "In 1895 the Linthwaite League joined from Lancashire and they employed a new coach, John Redfern, and was a South East professional."

Label: [1]

Sentence 1: "Bekhud Badayuni's most recent biographer was the late Dr. Asad Ahmad of Aligarh Muslim University's Urdu Department."

Sentence 2: "Budhayuni Muslim University's Urdu Department of the late Dr. Asad Ahmad Badek was Aligarh's most recent biographer."

Label: [1]

Sentence 1: "The party merged into the Deutsche Konservative Partei in 1946."

Sentence 2: "The Deutsche Partei merged into the Konservative party in 1946."

Label: [1]

**LABEL 2:** The two sentences have the same or similar meaning even though they are not exact copies of one another. Examples:

Sentence 1: "A keel-laying ceremony for the submarine was held at Electric Boat's Quonset Point facility in North Kingstown, Rhode Island, on 30 April 2007."

Sentence 2: "A keel-laying ceremony for the submarine was held on April 30, 2007 at Electric Boat's Quonset Point facility in North Kingstown, Rhode Island."

Label: [2]

Sentence 1: "In 1970, the remains of Sears, Smyth and Daly were repatriated to Ireland by The National Graves Association and given a military funeral with full honours."

Sentence 2: "In 1970, the remains of Sears, Daly and Smyth were repatriated to Ireland and given a full military funeral with honours by The National Graves Association."

Label: [2]

Sentence 1: "Since 2005, Maxim Vengerov has been Professor 1 at the Royal Academy of Music in London."

Sentence 2: "Since 2005, Professor Maxim Vengerov has been at the Royal Academy of Music in London."

Label: [2]

Following this order of going about your rating should help with efficiency:

1. First, just look at the second sentence and if it does not seem like a well-formed English sentence and appears to be meaningless and basically unreadable, rate it 0 and move on. The first sentence ("original") is a sentence from English wikipedia while the second one has no such guarantee of belonging to an English text.

2. If the second sentence seems okay when considered by itself, look at the first sentence and determine what you think should apply between 1 and 2, though 0 is still possible.

Figure 8: Instructions given to the annotators.