

# REFEREE: Reference-Free Sentence Summarization with Sharper Controllability through Symbolic Knowledge Distillation

Melanie Sclar<sup>1</sup> Peter West<sup>1</sup> Sachin Kumar<sup>2</sup> Yulia Tsvetkov<sup>1</sup> Yejin Choi<sup>1,3</sup>

<sup>1</sup>Paul G. Allen School of Computer Science & Engineering, University of Washington

<sup>2</sup>Language Technologies Institute, Carnegie Mellon University

<sup>3</sup>Allen Institute for Artificial Intelligence

msclar@cs.washington.edu

## Abstract

We present REFEREE, a novel framework for sentence summarization that can be trained reference-free (i.e., requiring no gold summaries for supervision), while allowing direct control for compression ratio. Our work is the first to demonstrate that reference-free, controlled sentence summarization is feasible via the conceptual framework of Symbolic Knowledge Distillation (West et al., 2022), where latent knowledge in pre-trained language models is distilled via explicit examples sampled from the teacher models, further purified with three types of filters: length, fidelity, and Information Bottleneck. Moreover, we uniquely propose *iterative* distillation of knowledge, where student models from the previous iteration of distillation serve as teacher models in the next iteration. Starting off from a relatively modest set of GPT-3-generated summaries, we demonstrate how iterative knowledge distillation can lead to considerably smaller, but better summarizers with sharper controllability. A useful by-product of this iterative distillation process is a high-quality dataset of sentence-summary pairs with varying degrees of compression ratios. Empirical results demonstrate that the final student models vastly outperform the much larger GPT3-Instruct model in terms of the controllability of compression ratios, without compromising the quality of resulting summarization.<sup>1</sup>

## 1 Introduction

We introduce REFEREE, a new framework for sentence summarization that works by iteratively generating and distilling knowledge into successively better models. This allows REFEREE to be [Re]ference fr[ee]—beginning by distilling from a large language model rather than with supervised data. Yet, our method results in a more efficient,

<sup>1</sup>See <https://github.com/msclar/referee> for code, models, and data.

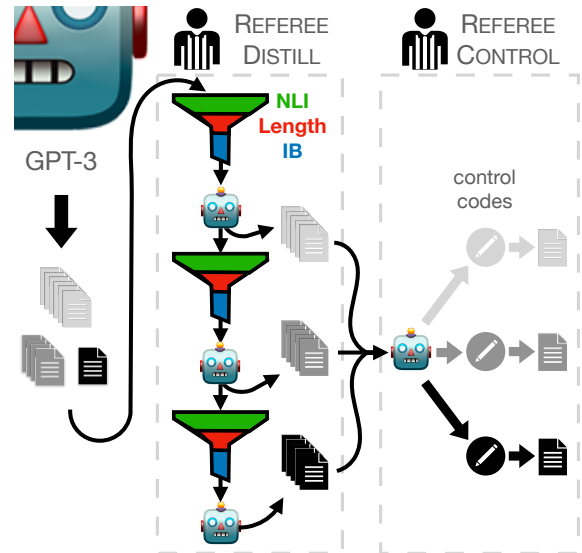


Figure 1: Our method results in high quality, *reference-free* compact summarizers. We begin by using a large language model (e.g. GPT-3) to generate many summaries that demonstrate different aspects we may want in a summary—**gray** represents an aspect well-represented in these generations, while **black** is under-represented. We first use REFEREE-DISTILL to iteratively filter and train summarizers that better represent these desirable aspects, e.g. shorter summary length. We then use generations from REFEREE-DISTILL to train a model in which these aspects are controllable: this is REFEREE-CONTROL.

compact, and controllable summarization model than what we start with.

Our work follows the paradigm of Symbolic Knowledge Distillation (West et al., 2022), which transfers implicit knowledge from a massive language model to a considerably smaller student model by explicitly generating knowledge in textual form. Unlike traditional knowledge distillation (Hinton et al., 2015) where the teacher model and the student model are of the same type, symbolic knowledge distillation allows for the student model to be of a different type.

Our work differs from West et al. (2022) in three

key aspects. First, our distillation is *iterative*: each student model becomes a teacher in successive rounds, refining and improving summarization at every step. Second, REFEREE controls for more than just overall quality, improving multiple model aspects in each round such as length, fidelity, and information bottleneck (Tishby et al., 1999), then allowing explicit length control at generation time. Third, our work is the first to show that reference-free, controlled sentence summarization can be formulated as symbolic knowledge distillation.

REFEREE works in two phases, illustrated in Figure 1. First, REFEREE-DISTILL uses a modest number of generated summaries from GPT-3 (Brown et al., 2020) to produce high quality and compact summarizers (Goyal et al., 2022). We follow an iterative approach; in each iteration we filter generations for desirable qualities, re-train a new and better summarizer, and finally generate new summaries for the next round. Each round amplifies effects of the previous rounds, improving notions of summary quality like entailment or shorter length. Second, REFEREE-CONTROL uses these iteratively distilled summaries to train a model with explicit control: in our experiments, we use progressively shortened generations from each iteration to train a final summarizer with explicit length control.

We find that REFEREE demonstrates compelling empirical results compared to competitive baselines. REFEREE-DISTILL, even without explicit length control, is able to generate shorter summaries with more consistency and equal quality compared with the original teacher model (GPT-3, 16x larger in size) as well as a supervised model. Moreover, REFEREE-CONTROL, which has more direct length control baked in, demonstrates a sharp degree of control in length, and succeeds at generating high quality summaries at specified lengths with significantly higher accuracy than GPT-3. In sum, the promising empirical results of REFEREE encourages further future investigation to extend the framework of symbolic knowledge distillation for reference-free, controlled text summarization.

## 2 Methods

We first describe REFEREE-DISTILL (see §2.1), an iterative procedure to promote specific behaviors that may not be prevalent in the original data, while maintaining summary quality. We explore two different filters, detailed in §2.2. We then de-

tail REFEREE-CONTROL (see §2.3), a model that separates summaries into categorical variables and is iteratively trained to, summarize a given sentence within the desired category (e.g., a range of compression ratio). In this work we only consider categories that reflect different compression ratio, but the same approach could be applied to other types of control categories, such as style.

### 2.1 Iterative Symbolic Knowledge Distillation: REFEREE-DISTILL

Let  $\mathcal{D} = \mathcal{D}_0 \cup \dots \cup \mathcal{D}_t$  denote a sentence corpus without reference summaries. We start with a teacher model (GPT3-Instruct Curie) from which we want to distill summarization knowledge under a fixed budget. Using  $\mathcal{D}_0$ —a small subset of  $\mathcal{D}$ —we first generate a dataset of sentence-summary pairs ( $\mathcal{C}_0$ ) by few-shot prompting the teacher and automatically filtering low-quality generations. Filters will be detailed in Section 2.2. Throughout the whole training procedure, we store each entry  $(s, s')$  as “ $s$  TL;DR:  $s'$  <eos>”. Here, <eos> denotes end of sequence and TL;DR; is a separator that has been shown to encourage summarization behavior (Radford et al., 2019).

Let  $\mathcal{M}_0$  be a pre-trained model significantly smaller than GPT-3 (GPT2-Large in our experiments). Using the seed dataset  $\mathcal{C}_0$ , we train a student model  $\mathcal{M}_1$  by fine-tuning  $\mathcal{M}_0$  with language modeling loss. We then iteratively refine this model by (1) using it to generate summaries for a subset of  $\mathcal{D}$ , (2) filtering them to remove undesired behaviors, and (3) training another student model on the filtered dataset, essentially distilling a better summarizer. More precisely,

$$\begin{aligned} \mathcal{C}_i &:= \text{filter}_i(\text{generate}(\mathcal{M}_i, \mathcal{D}_i)) \\ \mathcal{M}_{i+1} &:= \text{finetune}(\mathcal{M}_i, \mathcal{C}_i) \end{aligned}$$

We execute this procedure for  $t$  steps, creating  $t+1$  different summarization datasets in the process:  $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_t$ .<sup>2</sup> We discuss two possible instantiations of the  $\text{filter}_i$  below.

### 2.2 Filters

There is no one summary that is better than all others; depending on the desiderata of the end users, some might prefer shorter but less informative summaries, while others might prefer longer, and more informative ones. While some of these goals are

<sup>2</sup>Note that this process would stay identical if a user decided to use a human-generated summarization dataset as  $\mathcal{C}_0$ .

universal and always desired (for example, a summary should be accurate, in that it should not contain information not present in the input), others can be tailored to the end task. We use binary filters ( $\text{filter}_i$ ) to operationalize these goals. We experiment with the following filters.

**Summary Fidelity Filter** To encourage accurate summaries, we employ a simple but effective criterion: the summary should be entailed by the input sentence. More formally, we define a binary filter,  $f_{\text{NLI}}(s, s') := \mathbb{1}\{s \Rightarrow s'\}$ , and discard all non-entailed sentence-summary pairs to avoid using these samples when training the next iteration’s student. We measure entailment using an off-the-shelf state-of-the-art NLI model (Liu et al., 2022a).

**Summary Length Filter** While underexplored in prior work, constraining for the length of written text, especially in summarization, is a desirable feature to support real world applications with limited screen space. To obtain a corpus of summaries of varying lengths, at each distillation step  $i$ , we encourage the student  $\mathcal{M}_i$  to generate progressively shorter outputs. We achieve this by constraining  $\mathcal{C}_i$  to contain only summaries with a predefined compression ratio  $r_i \in [0, 1]$ . More precisely,

$$f_{\text{compress}}(s, s', r_i) = \mathbb{1}\left\{\frac{|s'|}{|s|} \leq r_i\right\}$$

where  $r_i > r_{i+1}$  for all  $i$ , to progressively summarize more succinctly.  $\frac{|s'|}{|s|}$  is commonly referred to as *compression ratio*. In theory, one could generate data for all desired compression ratios directly from  $\mathcal{M}_1$ . However, since the seed dataset  $\mathcal{C}_0$  is heavily skewed towards longer summaries, the final corpus after filtering with  $f_{\text{NLI}}$  would be extremely small for lower compression ratios. We find that combining the two filters and iteratively refining models to produce shorter, accurate summaries leads to a more diverse and still high-quality final corpus.

**Contextual Filter** For many applications, the sentences we need to summarize are part of a larger piece of text, such as a paragraph or a document (e.g. emails, articles). This contextual information may further improve sentence summary quality, since depending on the larger context, different information could be more important to be preserved, and inter-sentence redundancies could be removed. Inspired from West et al. (2019)’s interpretation of the Information Bottleneck principle (Tishby et al.,

1999), we consider the following filter:

$$f_{\text{NSP}} = \mathbb{1}\left\{\frac{p(s_{\text{next}}|s')}{p(s_{\text{next}}|s)} \geq l\right\}$$

where NSP refers to “next sentence prediction”,  $p$  is an oracle language model (which we approximate by GPT2-Large),  $s_{\text{next}}$  denotes the sentence immediately following the input sentence  $s$ , and  $l \in [0, 1]$  is a hyperparameter. Intuitively, we want to find summaries which are good predictors of the next sentence, to select the most crucial information and preserve coherence.  $l$  allows us to strike a balance between sacrificing some of the information in  $s$  and maintaining enough to predict  $s_{\text{next}}$ .

Adding  $f_{\text{NSP}}$  requires expanding the input sequence to also include the next sentence throughout the iterative distillation process defined in §2.1.

### Final REFEREE-DISTILL Filters Definition

We experiment with two filters,  $f_1$  and  $f_2$  (or #1 and #2, as we will refer to during experiments).  $f_1$  does not assume the existence of any context, and so it only filters for inaccuracies and length:

$$f_1(s, s'; s_{\text{next}}, r_i) = f_{\text{NLI}} \wedge f_{\text{compress}}$$

This allows  $f_1$  to be applied in broader contexts. We also define  $f_2$ , which adds contextual filtering:

$$f_2(s, s'; s_{\text{next}}, r_i, l) = f_{\text{NLI}} \wedge f_{\text{compress}} \wedge f_{\text{NSP}}$$

**Fluency Filter** To ensure fluency over several self-training iterations, we consider an additional filter only to be used in REFEREE-CONTROL. Given a sentence  $\mathbf{x} = (x_1, \dots, x_\ell)$ , we define  $\text{AvgNLL}(\mathbf{x}) := -\frac{1}{\ell} \sum_{i \leq \ell} \log p(x_i | x_{<i})$ . We determine a summary as fluent if and only if its mean Negative Log Likelihood (NLL) does not exceed that of source sentence, leading to the filter:

$$f_{\text{AvgNLL}}(s, s') = \mathbb{1}\left\{\text{AvgNLL}(s') \leq \text{AvgNLL}(s)\right\}$$

## 2.3 REFEREE-CONTROL

Using the high quality corpora of varying compression ratios obtained using REFEREE-DISTILL, we train REFEREE-CONTROL, a summarization model that allows explicit control for desired compression ratio. We divide all possible compression ratios into  $n$  buckets, where each bucket  $b_i = \left[\frac{i}{n}, \frac{i+1}{n}\right]$  for  $0 \leq i < n$ . Using  $b_i$  as control codes, we train a model that, when prompted with it, can summarize at a compression ratio within  $b_i$ .

Similar to  $\mathcal{D}$ , we start with a corpus  $\mathcal{F} = \mathcal{F}_0 \cup \dots \cup \mathcal{F}_t$  of sentences without reference summaries. Additionally, we create a seed corpus labeled with compression ratios,  $\mathcal{E}_0 = \mathcal{C}_0 \cup \dots \cup \mathcal{C}_t$  ( $\mathcal{F}_0 = \mathcal{D}_0 \cup \dots \cup \mathcal{D}_t$ ) now representing each example  $(s, s')$  as “ $s$  <sep> <bucket\_tok  $j$ > TL;DR:  $s'$  <eos>”, where <bucket\_tok  $j$ > corresponds to the bucket in which the example lies, that is  $\frac{|s'|}{|s|} \in b_j$ . <sep> is a special token. We denote each subset of  $\mathcal{E}_0$  corresponding to bucket  $j$  as  $\mathcal{E}_0^{(j)}$ . This seed dataset is filtered to remove low-quality generations, with the same filter as all the subsequent iterations.

Similar to REFEREE-DISTILL, starting with a pre-trained model  $\mathcal{N}_0$  (GPT2-Large), we train student models via iterative distillation. In each iteration  $i$ , (1) we fine-tune the student model using the bucket labeled corpus  $\mathcal{E}_i$ , (2) generate summaries for  $\mathcal{F}_i$  for all buckets, (3) filter them to create a new labeled corpus  $\mathcal{E}_{i+1}$ . We do not reinitialize the student at each iteration, but rather fine-tune starting from the teacher’s current local optima. We use  $h(s, s') = f_{\text{NLI}} \wedge f_{\text{AvgNLL}}$  as the filter. Formally,

$$\begin{aligned} \mathcal{N}_{i+1} &:= \text{finetune}(\mathcal{N}_i, \mathcal{E}_i^{(0)}, \mathcal{E}_i^{(1)}, \dots, \mathcal{E}_i^{(n-1)}) \\ \mathcal{E}_{i+1}^{(j)} &:= h(\text{generate}(\mathcal{N}_i, \mathcal{F}_i, j)) \quad \forall 0 \leq j < n \end{aligned}$$

## 2.4 Primal-Dual Problem Interpretation of Summarization

Assuming summaries are fluent and factual, sentence summaries trade off between two variables: level of compression and level of information preservation. We are able to effectively fix the level of compression by introducing control codes, and then develop models to maximize information preservation. This is our primal problem. Thanks to length-control codes, we can now also solve the dual problem: “what is the best shortest summary we could write?”. Written more precisely, given a fixed level of tolerance for losing information from the original sentence, what is the shortest summary we could write? Furthermore, comparing similar-lengthed summaries also allows for fairer comparisons, since we are effectively measuring changes in only one variable.

## 3 On GPT3’s Fidelity and Length Control

We analyze GPT3-Instruct Curie’s (Brown et al., 2020) sentence summarization capabilities. We promote GPT3 to summarize at different compression ratios by few-shot prompting with high-quality sentence-summary pairs in the desired compression

<i>Model</i>	<i>Avg. c.r. % (stdev)</i>	<i>WANLI Entailment %</i>	<i>c.r. <math>\geq 1</math> %</i>
GPT-3, c.r. 60-80%	82 (24)	<b>88</b>	31
GPT-3, c.r. 40-60%	78 (28)	81	33
GPT-3, c.r. 20-40%	<b>59 (28)</b>	76	<b>11</b>
Supervised baseline	55 (29)	71	7
Referee-Distill, filter #1	46 (15)	89	1
Referee-Distill, filter #2	49 (18)	91	2

Table 1: Statistics of automatically-generated datasets (ours and GPT-3). The first three rows refer to three different datasets generated through three-shot prompting GPT3-Instruct Curie with summaries from different compression ratios (c.r.). Following rows show results for the third and last iteration of our models, using each one of the two described filters ( $f_1, f_2$ , see §2.2). Sentences correspond to a held-out set during training.

ratios. More precisely, we do three-shot prompting with three different sets of summaries: one set of sentence-summary pairs has all three pairs with compression ratios in the interval  $[0.6, 0.8]$ , another set in  $[0.4, 0.6]$ , and another in  $[0.2, 0.4]$ .

We show that average compression ratio (c.r.) correlates with the prompts’ compression ratio (although variance is large), and up to 33% of the time models generate summaries longer than the original sentence (see Table 1). Qualitatively, this seems to be because of punctuation edits or hallucinations.

Besides using prompts that encourage shorter summaries, one can iteratively summarize through few-shot prompting. If  $f_p(s)$  is the summary GPT-3 generates when prompted with  $p$ , then  $f_p(f_p(\dots f_p(s))) = f_p^n(s)$  may also be a summary of  $s$ , possibly shorter. We find that successive application of the same prompt did not result in shorter summaries, i.e.  $|f_p^n(s)| \simeq |f_p(s)|$ , suggesting  $f_p$  is roughly idempotent in terms of length (see A.1).

These experiments motivate the need of more sophisticated approaches to length control and reliably summarizing without supervision.

## 4 Experiments

**Dataset** We create the corpora  $\mathcal{D}$  and  $\mathcal{F}$  by sampling contiguous sentence pairs from RealNews (Zellers et al., 2019) news articles. We filter out sentences shorter than 50 characters. Using GPT-3 as the teacher, we summarize sentences in  $\mathcal{D}_0$  and use the outputs with 60-80% compression ratio as our initial dataset  $\mathcal{C}_0$ , since it was the best one quantitatively and qualitatively. Although this implies that we will not initially have enough short sum-



Table 2: Results for paired comparison between two models (“comparative” and “baseline”, see first two columns). We use BERT Score ( $\in [0, 1]$ ), and ROUGE-1,2,L ( $\in [0, 100]$ ). Metric differences are computed as detailed in Section 4.1.1, and comparative model’ score is shown in brackets. A positive difference reflects an improvement using the comparative model over the baseline. The initial dataset consists of 10000 samples, from where we keep only summaries that differ in at most 10% compression ratio, and are not identical.

Comparative model	Baseline model	BERTScore F1 diff [ours]	R-1 diff [ours]	R-2 diff [ours]	R-L diff [ours]	% examples to compare	% equal summ.
Iter. 3, filter #1	GPT3, 60-80%	0.02 [0.77]	4.9 [59.1]	8.8 [47.8]	6.6 [58.1]	10%	2%
Iter. 3, filter #1	GPT3, 40-60%	0.03 [0.76]	5.4 [56.9]	5.4 [45.6]	7.2 [56.1]	14%	4%
Iter. 3, filter #1	GPT3, 20-40%	0.07 [0.74]	12.7 [53.9]	17.1 [42.4]	15.2 [53.0]	30%	4%
Iter. 3, filter #1	Supervised baseline	0.10 [0.74]	14.0 [53.8]	18.8 [42.4]	15.4 [53.0]	35%	0%
Iter. 1, filter #2	Iter. 1, filter #1	0.01 [0.83]	1.1 [69.5]	2.5 [58.4]	1.3 [68.0]	35%	26%
Iter. 2, filter #2	Iter. 2, filter #1	0.01 [0.80]	2.3 [63.3]	4.7 [52.9]	2.6 [62.3]	36%	27%
Iter. 3, filter #2	Iter. 3, filter #1	0.01 [0.75]	2.5 [55.4]	4.9 [45.3]	2.7 [54.7]	38%	26%

maries to train from, we will generate progressively shorter ones during our training procedure.

**Training Details** We use off-the-shelf model WANLI (Liu et al., 2022a) to create the filter  $f_{\text{NLI}}$ . We run 3 iterations of REFEREE-DISTILL with compression ratios for each training iteration are as follows:  $r_1 = 0.7$ ,  $r_2 = 0.5$ ,  $r_3 = 0.3$ . All generated data is decoded via beam search (with beam width 5). All experiments with  $f_{\text{NSP}}$  are done with  $r = e^{-6} \approx 0.0025$ , empirically decided through preliminary exploration. We train REFEREE-CONTROL with  $n = 10$  buckets for 7 iterations. For more details, refer to Appendix B.

While having contextual information for  $f_{\text{NSP}}$  is useful, many applications do not have contextual information available. Therefore, we use the corpus  $\mathcal{C}_0 \cup \dots \cup \mathcal{C}_t$  generated using  $f_1$  as initial training data for REFEREE-CONTROL with the goal of increasing its applicability.

**Supervised Baseline** We include a supervised baseline as a comparison point. Following Rush et al. (2015), we use Gigaword (Napoles et al., 2012; Graff et al., 2003) as a silver-labeled dataset where the headline is used as the summary for an article’s first sentence. We fine-tune GPT2-Large (Radford et al., 2019) (the same architecture as in our models) on this corpus, with default hyperparameters. Due to the training data’s nature, the supervised baseline has a low average compression ratio of 55% (similar to REFEREE-DISTILL Iteration 2, but longer on average than Iteration 3). Importantly, we do not (and cannot) include a conventional knowledge distillation from GPT-3 (e.g. Shleifer and Rush 2020) since the full distribution of token logits is unavailable.

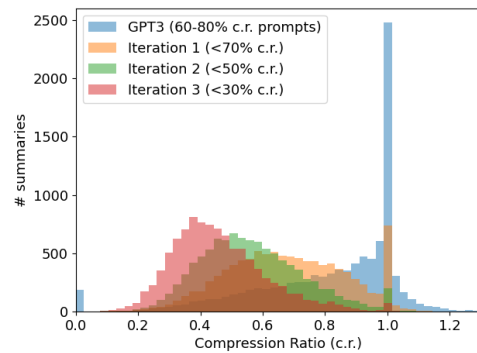


Figure 2: Histogram of compression ratios (c.r.) for  $\mathcal{C}_0, \dots, \mathcal{C}_3$  with filter #1 ( $f_1$ , see §2.2).  $\mathcal{C}_0$  are GPT-3 summaries, and  $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$  each of the training iterations.

## 4.1 Evaluating REFEREE-DISTILL

**Compression and Fidelity Statistics** We observe that iterative training with selection of progressively shorter summaries achieves the goal of generating shorter summaries, and additionally, less variance in compression ratios (Figure 2). Moreover, in Table 1 we observe that, by using an NLI filter during distillation, REFEREE-DISTILL summaries were  $\sim 90\%$  entailed by the original sentence according to WANLI (compared to 79% if not including an NLI filter during training). This vastly surpasses the comparable GPT-3 dataset (20-40%), and achieves similar fidelity as the best GPT-3 summaries, even when our model is significantly smaller. The same trends hold for both filters.

### 4.1.1 Comparison with GPT-3

We compare the quality of our summaries ( $s'$ ) and GPT-3 generated summaries ( $s''$ ), for every trained iteration and every GPT-3 dataset. Since

	<i>Faithful</i>	<i>Relevant</i>	<i>Fluent</i>
Supervised Baseline	0.778	0.883	0.838
GPT-3, 20-40%	0.825	0.950	0.935
<b>Referee-Distill Iter. 3 (us)</b>	<b>0.835</b>	<b>0.963</b>	<b>0.915</b>

Table 3: Results for human evaluation on 100 samples in faithfulness, relevance, and fluency, all in  $[0, 1]$  range (higher is better). See §4.1.2 for all details.

longer summaries will naturally be able to preserve more of the original information (Schluter, 2017), it is not reasonable to compare two wildly different compression ratios. Therefore, we only compare summaries that differ by length at most 10%:  $\left| \frac{|s'|}{|s|} - \frac{|s''|}{|s|} \right| \leq 0.1$ , where  $s$  is the original sentence.

To measure summary quality automatically, we compute BERTScore (Zhang\* et al., 2020) and ROUGE-1,2,L (Lin, 2004) against the original sentence  $s$  (no references are available). Given a metric  $m$ , we evaluate models based on  $m(s', s) - m(s'', s)$ . Positive values reflect that  $s'$  had higher scores than the baseline summary  $s''$ , which is desirable for all our metrics. Our models show significant improvements in all metrics when compared to every GPT-3 dataset, and the supervised baseline (see Table 2). Our model shows especially large improvements when compared to the shortest GPT-3 summaries (20-40% prompts). This suggests our iterative procedure was able to preserve quality better during the selection for shorter summaries.

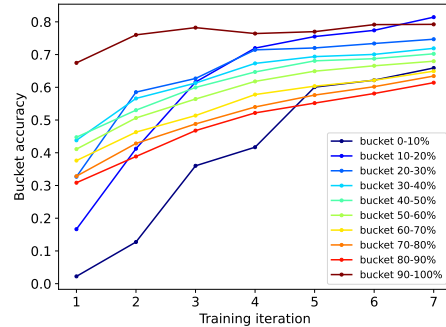
Finally, we compare the effect of introducing a contextual filter  $f_{\text{NSP}}$ . We compare two identical training runs that differ only in the filter applied (filter #2 vs. #1), and we observe small improvements when including  $f_{\text{NSP}}$  (see last rows of Table 2, higher values mean #2 was better than #1).

#### 4.1.2 Human Evaluation

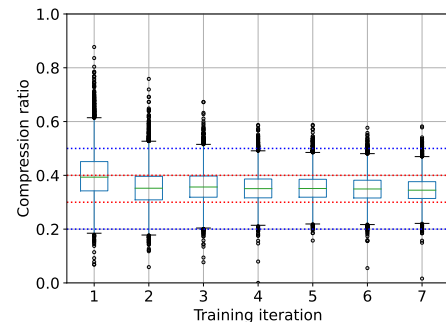
We conduct a human evaluation to verify the qualities of summaries from REFEREE-DISTILL. We measure 3 axes: faithfulness (*is the summary true to the source?*), relevance (*does the summary capture important information from the source?*) and fluency, each on a 3-point Likert scale.<sup>3</sup> We conduct our evaluation on 100 examples and find agreement by Fleiss  $\kappa$  (Fleiss, 1971) of 0.32, 0.34, and 0.57 (respectively) indicating fair to moderate agreement (Landis and Koch, 1977).

We compare between different methods to obtain succinct summaries: REFEREE-DISTILL (Iteration

<sup>3</sup>Note that humans did not agree well between the top two buckets for relevance, so we combined these into a single bucket and calculate average and agreement on this.



(a) REFEREE-CONTROL bucket accuracy per iteration.



(b) Compression ratio boxplot per iteration for bucket 30-40%. Red lines represent bucket bounds, and blue lines represent a 10% compression ratio difference to bucket bounds. Other buckets shown in Appendix C.2.

Figure 3: Statistics per REFEREE-CONTROL iteration, computed on the datasets  $\mathcal{E}_i$  generated during training.

3, function #1), GPT-3 20-40%, and the supervised baseline. Following §4.1.1, we only compare sentences if all generations differ in compression ratio of at most 10%. See results in Table 3.

Broadly, we find that REFEREE-DISTILL and GPT-3 achieve significantly higher quality than the supervised baseline, with REFEREE-DISTILL showing even slightly higher scores than GPT-3 for all 3 axes. We also note that this evaluation may somewhat favor baselines, since we only select examples where they achieved a similar compression ratio as us. We are not accounting for the fact that REFEREE-DISTILL can generate short summaries for many examples the other two systems cannot.

#### 4.2 Evaluating REFEREE-CONTROL

We train all our models with  $n = 10$  buckets to provide a very fine-grained control: the average sentence length in our dataset is 134 characters long, which implies that each bucket may span  $\sim 13$  characters for the average sentence. This implies a model may only have one or two words of freedom

Model	BERT Score	R-1	R-2	R-L	Bucket Acc.
GPT3, 60-80%	0.66	40.0	24.3	36.8	3%
GPT3, 40-60%	0.65	39.5	28.0	36.9	5%
GPT3, 20-40%	0.63	34.1	18.4	31.0	14%
Referee-Ctrl, Iter. 3	<b>0.69</b>	47.7	<b>36.7</b>	47.3	61%
Referee-Ctrl, Iter. 5	0.68	46.5	34.8	<b>46.0</b>	69%
Referee-Ctrl, Iter. 7	0.66	44.5	32.1	43.9	<b>71%</b>

Table 4: BERT Score, ROUGE-1,2,L and bucket accuracy for the 30-40% bucket. Three GPT3 datasets are shown, and three different iterations of REFEREE-CONTROL. Data for other buckets is in Appendix C.1.

before exceeding the maximum length allowed in a bucket. It is important to note that models do not have direct access to a mapping of subword token to character length, and rather need to estimate character length during training. We use character length constrains since text applications will impose this type of restriction.

We show that each REFEREE-CONTROL iteration increases bucket accuracy and reduces resulting compression rate variance (See Figure 3, Appendix C.2). To maximize quality, in all our experiments we only use one sampled beam. If we wished to maximize bucket accuracy at the expense of possibly reduced quality, we can take the top beams and select the most likely one that is in the prompted bucket. This procedure increases the bucket accuracy dramatically: in iteration two, using one beam has a 42% bucket accuracy for the bucket 80-90% (in a held-out set), whereas using three yields an accuracy of 71%, and five, 82%. This same trend holds for other iterations and buckets, reaching 93% bucket accuracy in iteration 7.

This trade-off between bucket accuracy and summary quality that can be seen for bucket  $b_3$  in Table 4, although the behavior is consistent for all buckets (See Appendix C.1). There, Iteration 3 has slightly higher BERTScore and ROUGE than Iteration 7, at the expense of lower bucket accuracy. We believe this is because as bucket accuracy increases we reach harder to summarize examples at the desired length range, causing average scores to drop.

Later iterations also show more (small) disfluencies, which we partially attribute to the aforementioned cause. Also, small disfluencies may propagate over time, which we mitigate by using the mean negative log likelihood ratio filter  $f_{\text{AvgNLL}}$ . Removing the fluency filter  $f_{\text{AvgNLL}}$  will also enable 9 more points of bucket accuracy on average.

Method	$acc$	$acc^+$
<b>High compression (40-60%)</b>		
REFEREE-CONTROL Iter. 3	0.768	0.360
GPT-3 40-60%	0.194	0.095
Ghalandari et al. (2022)	0.530	0.153
Liu et al. (2022b)	0.320	0.043
Schumann et al. (2020)	0.357	0.037
<b>Extreme compression (20-40%)</b>		
REFEREE-CONTROL Iter. 3	0.670	0.233
GPT-3 20-40%	0.250	0.117
Ghalandari et al. (2022)	0.431	0.081
Liu et al. (2022b)	0.333	0.034
Schumann et al. (2020)	0.337	0.048

Table 5: Accuracy of different models at generating high-quality, length-controlled summaries for two ranges, 20-40% and 40-60%. Summaries must fulfill length constraints and meet human notions of quality at reasonable ( $acc$ ) or high ( $acc^+$ ) standards to be accepted. Higher is better. See setup details in C.3.

**Importance of Iterative Distillation** We trained a GPT2-Large model for 10 epochs with the same seed dataset  $\mathcal{E}_0$ , and compare with REFEREE-CONTROL iteration 5 (also trained for 10 epochs in total). Bucket accuracy of REFEREE-CONTROL was  $\sim 20$  points better: for 30-40% bucket, bucket accuracy in the non-iterative version was 42% vs. 69% for REFEREE-CONTROL; for 70-80% bucket, accuracy was 34% vs. 58%.

#### 4.2.1 Human Evaluation

We aim to explicitly test the capacity of systems to generate an *acceptable summary*. That is, a summary to meets minimum human measures of quality, and also adheres to the desired length constraint. We omit the supervised baseline here, as it does not have explicit length control, and thus include only REFEREE-CONTROL and GPT-3.

Specifically, we measure summary accuracy as the fraction of summaries (of 100 randomly pre-selected sentences) that adhere to length control while being sufficiently fluent, relevant, and faithful. These axes are measured as in §4.1.2, achieving agreement by Fleiss  $\kappa$  (Fleiss, 1971) of 0.34, 0.22, and 0.25 (respectively) indicating fair agreement (Landis and Koch, 1977). We include two accuracy measurements:  $acc$ , requiring adhering to length constraints as well as at least 2 (“fair”) out of 3 on all human measures of quality; and  $acc^+$ , which requires 3 out of 3 on all measures along with length adherence.

Table 5 includes results for the 20-40% and

40-60% compression ranges, following our GPT-3 datasets (see C.3 for more setup details). REFEREE-CONTROL vastly outperforms GPT3 for both regimes and metrics. More precisely, for the 40-60% regime REFEREE-CONTROL showed +296% in *acc* and +279% in *acc*<sup>+</sup> when compared with GPT3; and for the 20-40% regime, REFEREE-CONTROL obtained +68% and +99% respectively. We additionally included three unsupervised summarization systems in the human evaluation. These systems perform some length control, making their summaries comparable, but all three models performed poorly when compared with REFEREE-CONTROL: they were at least 23 points below REFEREE-CONTROL in *acc*, and 17 points below in *acc*<sup>+</sup>.

Lastly, we would like to emphasize that REFEREE-CONTROL aims to summarize all examples at the requested compression, regardless of the original sentence’s length or difficulty. GPT3, on the other hand, only summarizes in the requested compression for longer sentences, which generally correlate with easier cases (see details in A.3).

## 5 Related Work

**Unsupervised Summarization** The vast majority of prior work in sentence summarization assumed access to large-scale text-summary paired datasets from which to train supervised models (Rush et al., 2015; Nallapati et al., 2016; Narayan et al., 2018). Nonetheless, these datasets are costly to create, and naturally-occurring summarization datasets (such as news highlights) are noisy and not easily found in other domains. Therefore, recent work emphasized the need for developing unsupervised or self-supervised methods such as autoencoders (Miao and Blunsom, 2016; Baziotis et al., 2019), but in general they lead to less fluent summaries; more recent work has explored the Information Bottleneck Principle (West et al., 2019) instead of the reconstruction loss of autoencoders. Our work contributes to this emerging line of research by demonstrating an entirely different method based on symbolic knowledge distillation.

**Length-Controlled Summarization** While real world applications would require controlling for summary length, most prior work for automatic summarization has not proposed a principled mechanism for controlling the level of compression. Notable exceptions include Kikuchi et al. (2016) and He et al. (2020); Fan et al. (2018); Liu et al. (2018).

These last works developed supervised models for controllable length summarization by adding control codes that corresponded to a range of summary lengths—commonly referred to as *buckets*. However, in both works the degree of control is heavily dependent on the training dataset, since bucket bounds are defined so that each one has the same number of examples; this may make one bucket correspond to a wide range of compression ratios. Our work adds a unique contribution by proposing a reference-free method that allows for full range of controls, and explicitly evaluates for that behavior.

Concurrently to this work, Ghalandari et al. (2022) and Liu et al. (2022b) (the latter optimizing Schumann et al., 2020) proposed unsupervised mechanisms that enforce length at the word level, either through a reward mechanism or with strict length enforcement. In contrast, our method uses iterative knowledge distillation to achieve length control. We also control length at the character level, rather than word level. This can result in a more fair and challenging notion of control, as it prevents performing the simple strategy in which all function words are removed first to maximize general meaning. Notably, and also in contrast to our work, both Ghalandari et al. (2022) and Liu et al. (2022b) require training separate models for compressing at different compression ratios.

**Knowledge Distillation** Many prior works have focused on similar notions of transferring knowledge between models through generation and distillation, and we draw particular inspiration from West et al. (2022). Shleifer and Rush (2020) also follow a similar form to our work, distilling a summarizer from pretrained models. Our work differs in two key ways. First, like most distillation works, Shleifer and Rush (2020) assume having a model trained for the task and with access to its full distribution of token logits—both not the case of GPT-3, used here. Second, like many distillation studies, Shleifer and Rush (2020) aim to mimic the teacher model’s distribution, while we attempt to improve it. This core detail sets us apart from many works employing a large teacher model (Kim and Rush, 2016; Schick and Schütze, 2021; Ye et al., 2022), teaching a student to mimic a distribution rather than improve it as in our case.

**Natural Language Inference (NLI) for Summarization** Pasunuru et al. (2017); Pasunuru and



Bansal (2018); Li et al. (2018) have used NLI for summarization enhancement: Pasunuru et al. (2017) use entailment in multi-task learning, and Pasunuru and Bansal (2018); Li et al. (2018) use entailment probability as a reward. In this work, we propose an alternative approach for incorporating NLI for enhancing fidelity of summarization under the Symbolic Knowledge Distillation framework.

## 6 Conclusions

We presented REFEREE, a framework for sentence summarization that can be trained without reference summaries, while allowing direct control for summary compression ratio. We uniquely proposed *iterative* Symbolic Knowledge Distillation, where student models from the previous iteration of distillation serve as teacher models in the next. Distilled models are significantly smaller than the original teacher, GPT-3, and empirical results demonstrated that the final student models vastly outperform the much larger GPT3-Instruct model in terms of the controllability of compression ratios, without compromising the resulting summaries' quality. A useful by-product of this iterative distillation process is a high-quality sentence summarization dataset with varying degrees of compression, which we will release jointly with our models upon publication.

## Acknowledgements

We thank Vidhisha Balachandran, Alisa Liu, and the whole Tsvetshop lab for the valuable discussions. We also thank OpenAI for providing academic access to the GPT-3 API. This material is based upon work partly funded by the DARPA CMO under Contract No. HR001120C0124, and by DARPA MCS program through NIWC Pacific (N66001-19-2-4031). S.K. has been supported by a Google PhD Fellowship. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily state or reflect those of the United States Government or any agency thereof.

## Limitations

REFEREE was entirely developed and tested with sentences extracted from news articles. More work is needed to assess REFEREE's robustness when applying it to other domains. These domains may differ in text type, topic, or even temporal differences, that may cause a distribution shift.

REFEREE's success is also tied to other systems' quality, mainly the seed dataset generator (GPT-3) and the summary fidelity filter (operationalized using WANLI entailments). REFEREE may propagate errors and biases in NLI entailments, which may be remedied in the future as NLI research progresses. We believe some edge cases generated by REFEREE may be useful to further augment data in NLI systems, but that investigation was outside the scope of this paper.

REFEREE is built entirely at the sentence level, and more work is needed to extend it to paragraph or document-level, although some of the same ideas could be applied (e.g., control codes over longer inputs).

## References

- Christos Baziotis, Ion Androutsopoulos, Ioannis Konstantas, and Alexandros Potamianos. 2019. [SEQ<sup>3</sup>: Differentiable sequence-to-sequence-to-sequence autoencoder for unsupervised abstractive sentence compression](#). 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 673–681, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Angela Fan, David Grangier, and Michael Auli. 2018. [Controllable abstractive summarization](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Demian Gholipour Ghalandari, Chris Hokamp, and Georgiana Ifrim. 2022. Efficient unsupervised sentence compression by fine-tuning transformers with reinforcement learning. *arXiv preprint arXiv:2205.08221*.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.

- Junxian He, Wojciech Kryściński, Bryan McCann, Nazneen Rajani, and Caiming Xiong. 2020. Ctrlsum: Towards generic controllable text summarization. *arXiv preprint arXiv:2012.04281*.
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. 2015. Distilling the knowledge in a neural network. *NIPS Deep Learning and Representation Learning Workshop*, 2(7).
- Yuta Kikuchi, Graham Neubig, Ryohei Sasano, Hiroya Takamura, and Manabu Okumura. 2016. Controlling output length in neural encoder-decoders. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1328–1338.
- Yoon Kim and Alexander M. Rush. 2016. **Sequence-level knowledge distillation**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Haoran Li, Junnan Zhu, Jiajun Zhang, and Chengqing Zong. 2018. Ensure the correctness of the summary: Incorporate entailment knowledge into abstractive sentence summarization. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1430–1441.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Alisa Liu, Swabha Swayamdipta, Noah A Smith, and Yejin Choi. 2022a. Wanli: Worker and ai collaboration for natural language inference dataset creation. *arXiv preprint arXiv:2201.05955*.
- Puyuan Liu, Chenyang Huang, and Lili Mou. 2022b. **Learning non-autoregressive models from search for unsupervised sentence summarization**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7916–7929, Dublin, Ireland. Association for Computational Linguistics.
- Yizhu Liu, Zhiyi Luo, and Kenny Zhu. 2018. Controlling length in abstractive summarization using a convolutional neural network. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4110–4119.
- Yishu Miao and Phil Blunsom. 2016. Language as a latent variable: Discrete generative models for sentence compression. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 319–328.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. **Abstractive text summarization using sequence-to-sequence RNNs and beyond**. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Courtney Napoles, Matthew Gormley, and Benjamin Durme. 2012. Annotated gigaword. pages 95–100.
- 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. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807.
- Ramakanth Pasunuru and Mohit Bansal. 2018. **Multi-reward reinforced summarization with saliency and entailment**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 646–653, New Orleans, Louisiana. Association for Computational Linguistics.
- Ramakanth Pasunuru, Han Guo, and Mohit Bansal. 2017. Towards improving abstractive summarization via entailment generation. In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 27–32.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. **A neural attention model for abstractive sentence summarization**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389, Lisbon, Portugal. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. **Generating datasets with pretrained language models**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6943–6951, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Natalie Schluter. 2017. The limits of automatic summarisation according to rouge. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 41–45.
- Raphael Schumann, Lili Mou, Yao Lu, Olga Vechtomova, and Katja Markert. 2020. **Discrete optimization for unsupervised sentence summarization with word-level extraction**. pages 5032–5042.
- Sam Shleifer and Alexander M Rush. 2020. Pre-trained summarization distillation. *arXiv preprint arXiv:2010.13002*.

- Naftali Tishby, Fernando C. Pereira, and William Bialek. 1999. The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, pages 368–377.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena D Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. Symbolic knowledge distillation: from general language models to commonsense models. *To appear in NAACL-HLT*.
- Peter West, Ari Holtzman, Jan Buys, and Yejin Choi. 2019. [BottleSum: Unsupervised and self-supervised sentence summarization using the information bottleneck principle](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3752–3761, Hong Kong, China. Association for Computational Linguistics.
- Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022. Zerogen: Efficient zero-shot learning via dataset generation. *ArXiv*, abs/2202.07922.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. Defending against neural fake news. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 9054–9065. Curran Associates, Inc.
- Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

## A Further Insights on GPT-3 Datasets

### A.1 Successive Application of GPT-3 Prompts




Figure 4: Histogram showing compression ratios of successive application of the same summarization prompts. Graphic shows no real difference in repeating the same prompt using the previously generated summary to summarize further.

### A.2 Compression Ratio Distribution per Dataset




Figure 5: Histogram of compression ratio distribution for each prompt set.

### A.3 Distribution of Sentence Length in Cases Where GPT-3 Respected the Desired Compression

We observe that the distribution of original sentence length is markedly different when considering all the sentences in the dataset, versus when only considering the sentences where GPT-3’s summary was in the desired compression range (see Figure 6. This shows that GPT-3’s sentences compresses

more only in easier cases, whereas REFEREE-CONTROL follows the distribution of the original set (see 7).

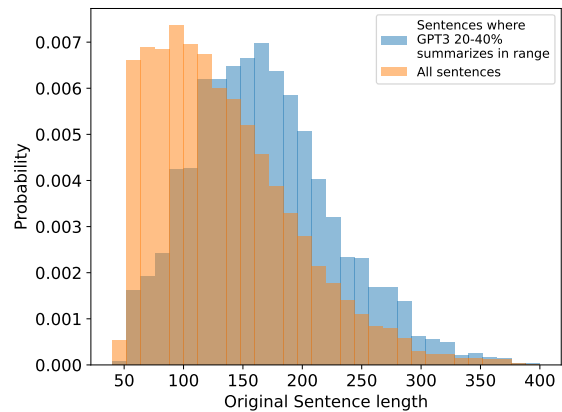


Figure 6: Distribution of original sentence length of all original sentences, and of the subset of samples where GPT-3 20-40% returned a summary in the prompted compression range.

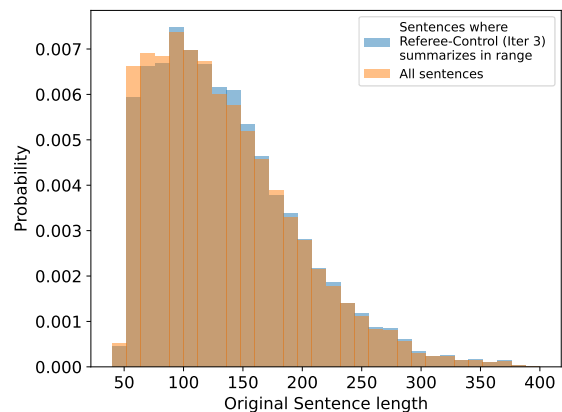


Figure 7: Distribution of original sentence length of all original sentences, and of the subset of samples where REFEREE-CONTROL 30-40% returned a summary in the 20-40% compression range (following the human evaluation setup).

## B Training and Dataset Details

We use GPT2-Large for all our fine-tuned models (774M weights, ~16x smaller than GPT3-Curie (Radford et al., 2019; Brown et al., 2020)), and fine-tune for 5 epochs during each iteration of REFEREE-DISTILL, and for 2 epochs for REFEREE-CONTROL.  $\mathcal{D}_0$  consists of 100000 sentences,  $\mathcal{D}_{i>0}$  consists of 40000 sentences each;  $\mathcal{F}_0$  consists of 220000 sentences,  $\mathcal{F}_{i>0}$  consists of 10000 sentences each. All generated data is decoded with sampling beam search with 5 beams.



We balance each bucket before training each REFEREE-CONTROL iteration to avoid overrepresenting some classes. We use string characters 0 to  $n - 1$  as control codes for their respective buckets. We find that repeating the control code increases bucket accuracy –likely because the model attends more to these tokens–, so the end control code for bucket 2 will be 2 2 2 2 2 2 2 2 2 2 (ten repetitions).

## C Additional Details for REFEREE-CONTROL

### C.1 Performance for Additional Buckets

Model	BERT Score	R-1	R-2	R-L	Bucket Acc.
GPT3, 60-80%	0.76	57.5	41.9	54.6	7%
GPT3, 40-60%	0.76	58.9	48.1	56.6	10%
GPT3, 20-40%	0.71	48.9	32.1	44.8	13%
Referee-Ctrl, Iter. 3	0.79	65.1	54.9	64.5	56%
Referee-Ctrl, Iter. 5	0.79	64.7	54.1	64.2	66%
Referee-Ctrl, Iter. 7	0.78	63.2	51.9	62.7	68%


Table 6: BERT Score, ROUGE-1,2,L and bucket accuracy for the 50-60% bucket. Three GPT3 datasets are shown, and three different iterations of REFEREE-CONTROL.

Model	BERT Score	R-1	R-2	R-L	Bucket Acc.
GPT3, 60-80%	0.84	72.3	58.1	70.0	11%
GPT3, 40-60%	0.83	71.6	62.4	69.3	10%
GPT3, 20-40%	0.79	62.4	46.9	58.3	8%
Referee-Ctrl, Iter. 3	0.88	79.3	71.4	78.7	48%
Referee-Ctrl, Iter. 5	0.88	78.9	70.9	78.5	57%
Referee-Ctrl, Iter. 7	0.87	77.6	68.9	77.3	63%


Table 7: BERT Score, ROUGE-1,2,L and bucket accuracy for the 70-80% bucket. Three GPT3 datasets are shown, and three different iterations of REFEREE-CONTROL.

It is crucial to note that because variance reduces with the number of iterations, cases where the summary did not land on the requested compression ratio are close to fulfilling the constraint. For example, when prompting with 30-40%, 87.5% of the samples fell in the 20-40% range for REFEREE-CONTROL Iteration 7 (77% for Iteration 3; 83% for Iteration 5).


### C.2 REFEREE-CONTROL Compression Ratio Distribution per Training Iteration



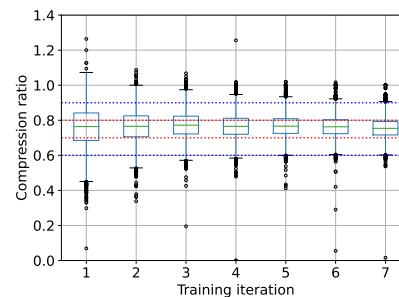
(a) Bucket 10-20%



(b) Bucket 40-50%



(c) Bucket 50-60%



(d) Bucket 70-80%

Figure 8: Boxplots showing the compression ratio distribution per training iteration, per bucket. 30-40% bucket is shown in Figure 3 in the main paper.

### C.3 Human Evaluation Setup Details

For the 20-40% range, we compare GPT-3 against REFEREE-CONTROL conditioned for 30-40%, and accept all summaries within the same range as GPT-3 prompts (20-40%). Similarly, for 40-60% we condition REFEREE-CONTROL with 50-60%, and accept all summaries within 40-60% to equalize comparisons.

Regarding baselines, we trained Ghalandari et al. (2022) from scratch using RealNews. Schumann et al. (2020) was used out-of-the-box, and we use the default training data for Liu et al. (2022b) –the publicly available Schumann et al. (2020) generations– due to time constraints. All baselines measure compression ratio at the word level, in contrast to our work that does it at the character level. We train with a 30% target compression ratio in baselines when comparing with the 20-40% regime, and 50% when comparing with the 40-60% regime. This is done to maximize the number of examples that fall in the desired range, since some baselines enforce target compression ratio more strictly than others.

### D On BERT Score Monotonicity

In §2.4, we discussed that introducing control codes allows us to solve a related question: given we have a specific tolerance for information loss  $k$ , what is the shortest summary we could write? Estimating the level of information preservation with BERT Score Recall, we could sample from each bucket and select the shortest summary with BERT Score Recall  $\geq k$ . BERT Score Recall scores are generally well-ordered when ordering increasingly by bucket: we see that on average, the longest non-decreasing subsequence of BERT Score Recall scores is  $\sim 7.7$ - $7.8$  for all iteration steps. Having a non-decreasing subsequence of 10 would mean that BERT Scores are perfectly ordered (a summary in a higher bucket would have always higher BERT Score recall than a lower bucket). 7.8 means that scores are generally well-ordered, with some noisy generations.

### E Examples of Generations

#### E.1 Random sample of generations used in the 40-60% human evaluation

*Original:* Viggo left South America aged 11, when his parents divorced and he moved to upstate New York with his mother and siblings.

*GPT3:* Viggo left South America when his parents divorced. **[41.8%, in range]**

*REFEREE-CONTROL:* Viggo left South America aged 11, when his parents divorced him. **[52.46%, in range]**

*Ghalandari et al. (2022):* Viggo left South America aged 11 when his parents divorced. **[48.36%, in range]**

*Liu et al. (2022b):* parents divorced moved to upstate new york with his mother and siblings **[58.2%, in range]**

*Original:* And Yahoo appears to have had another money-losing fumble on its hands with the free livestream of the NFL's Oct. 25 Bills-Jaguars game from London.

*GPT3:* Yahoo appears to have had another money-losing fumble on its hands with the free livestream of the NFL's Oct. 25 Bills-Jaguars game from London. **[97.3%, 37.3% too long]**

*REFEREE-CONTROL:* Yahoo appears to have had another money-losing fumble on its hands with the free livestream. **[62.16%, 2.16% too long]**

*Ghalandari et al. (2022):* And Yahoo appears to have had another money-losing fumble with livestream Oct. 25 Bills-Jaguars. **[64.86%, 4.86% too long]**

*Liu et al. (2022b):* and yahoo have hands of nfl game **[21.62%, 18.38% too short]**

*Original:* Suárez almost levelled with his latest sweeping free-kick from distance but the clearer chances came to Hull against a ponderous Liverpool defence.

*GPT3:* Suarez almost levelled with his latest sweeping free-kick from distance but Hull's chances came against a ponderous Liverpool defense. **[91.16%, 31.16% too long]**

*REFEREE-CONTROL:* Suarez almost levelled with his latest sweeping free-kick from distance. **[48.98%, in range]**

*Ghalandari et al. (2022):* Suárez almost levelled with his free-kick to Hull Liverpool. **[40.82%, in range]**

*Liu et al. (2022b):* <lunk> from distance the clearer chances came to hull liverpool <lunk> **[48.98%, in range]**

*Original:* Specifically, the company agreed to limit pre-payments, to provide accurate estimates of charges, and to disclose details of financing

agreements.

*GPT3*: Verizon has agreed to limit pre-payments, to provide accurate estimates of charges, and to disclose details of financing agreements. [90.41%, 30.41% too long]

REFEREE-CONTROL: The company agreed to limit pre-payments and to provide accurate estimates of charges. [58.9%, in range]

*Ghalandari et al. (2022)*: Specifically, the company agreed to limit pre-payments to estimates financing. [53.42%, in range]

*Liu et al. (2022b)*: company agreed to limit to provide accurate estimates of disclose details financing [56.85%, in range]

*Original*: "There was a family in the car, they got out before it hit.

*GPT3*: There was a family in the car when the car hit a tree, and they all got out before it hit the tree. [167.8%, 107.8% too long]

REFEREE-CONTROL: There was a family in the car. [50.85%, in range]

*Ghalandari et al. (2022)*: "There was a family in the car. [52.54%, in range]

*Liu et al. (2022b)*: there was a family they got before it [62.71%, 2.71% too long]

*Original*: As for the coming revolution, I can't say for certain when that will happen.

*GPT3*: I can't say for certain when the coming revolution will happen. [82.89%, 22.89% too long]

REFEREE-CONTROL: I can't say for certain when it will happen. [57.89%, in range]

*Ghalandari et al. (2022)*: As for the coming revolution I can't say. [53.95%, in range]

*Liu et al. (2022b)*: for the coming revolution i can <lunkl> t say [59.21%, in range]

*Original*: In one New Jersey county, the Red Cross was AWOL, says Lt. Matthew Tiedemann of the Bergen County Office of Emergency Management.

*GPT3*: In one New Jersey county, the Red Cross was AWOL. [37.98%, 2.02% too short]

REFEREE-CONTROL: In one New Jersey county, the Red Cross was AWOL, says Lt. Matthew Tiedemann. [59.69%, in range]

*Ghalandari et al. (2022)*: In one New Jersey county Red Cross was AWOL says Lt. Matthew Tiedemann. [55.04%, in range]

*Liu et al. (2022b)*: in new jersey county red cross

awol lt. <lunkl> emergency management [52.71%, in range]

*Original*: Furthermore, though some might argue that the purpose of food is merely to remind us of our own mortality, I would argue that the best cuisine can be an art form.

*GPT3*: Food can be both utilitarian and artful. [24.69%, 15.31% too short]

REFEREE-CONTROL: Furthermore, I would argue that the best cuisine can be an art form, rather than a food source. [58.64%, in range]

*Ghalandari et al. (2022)*: Furthermore of food mortality I would argue that the best cuisine can be an art form. [52.47%, in range]

*Liu et al. (2022b)*: the is our would that the best cuisine can be an art [32.1%, 7.9% too short]

*Original*: Without proper conservation facilities nearby, the artifacts Batchvarov and his colleagues discovered had to be reburied underwater.

*GPT3*: Without proper conservation facilities nearby, the artifacts Batchvarov and his colleagues discovered had to be reburied underwater. [100.0%, 40.0% too long]

REFEREE-CONTROL: The artifacts Batchvarov and his colleagues discovered had to be reburied. [56.06%, in range]

*Ghalandari et al. (2022)*: Without conservation artifacts Batchvarov had to reburied underwater [51.52%, in range]

*Liu et al. (2022b)*: conservation facilities artifacts colleagues discovered to be reburied underwater [61.36%, 1.36% too long]

*Original*: Often people are desperate or ridiculous rather than heroic, so I do what needs to be done."

*GPT3*: That person is heroic. [23.91%, 16.09% too short]

REFEREE-CONTROL: Often people are desperate or ridiculous, so I do it. [57.61%, in range]

*Ghalandari et al. (2022)*: Often people are desperate or ridiculous heroic I. [54.35%, in range]

*Liu et al. (2022b)*: so i do what needs to be done <lunkl> [40.22%, in range]

*Original*: Liam's feet make a guest appearance in the picture, as does his Great Dane Watson who he bought with ex-girlfriend Sophia Smith can be seen by his settee.

*GPT3*: Liam's feet and his Great Dane Watson can be seen by his settee in the picture. [51.3%, in range]

REFEREE-CONTROL: Liam's feet make a guest appearance in the picture, as does his Great Dane Watson. [53.25%, in range]

*Ghalandari et al. (2022)*: Liam's feet make a guest appearance in the picture as Great Dane Watson ex-girlfriend settee. [60.39%, 0.39% too long]

*Liu et al. (2022b)*: <lunk> feet picture does dane watson he bought with ex-girlfriend sophia smith be <lunk> [58.44%, in range]

*Original*: I like this little gallery, the Small Weston Room, where the paintings go right up to the ceiling.

*GPT3*: I like this little gallery, the Small Weston Room. [51.02%, in range]

REFEREE-CONTROL: I like this little gallery, the Small Weston Room. [51.02%, in range]

*Ghalandari et al. (2022)*: I like this little gallery Small Weston Room where paintings ceiling. [70.41%, 10.41% too long]

*Liu et al. (2022b)*: i like room where the go right to the ceiling [45.92%, in range]

*Original*: After a stint as a lobbyist, during which he helped establish two think tanks affiliated with James Dobson's hugely influential Focus on the Family, Parker became a confidant of Roy Moore, then a county judge.

*GPT3*: After a stint as a lobbyist, during which he helped establish two think tanks affiliated with James Dobson's hugely influential Focus on the Family, Parker became a confidant of Roy Moore, then a county judge. [100.0%, 40.0% too long]

REFEREE-CONTROL: After a stint as a lobbyist, Parker became a confidant of Roy Moore, then a county judge, during a stint. [50.24%, in range]

*Ghalandari et al. (2022)*: After a stint as a lobbyist helped think tanks James Dobson, Parker became a confidant of Roy Moore. [47.85%, in range]

*Liu et al. (2022b)*: a a lobbyist s focus on the family parker became confidant roy moore [32.54%, 7.46% too short]

*Original*: Germany will be the favourites but it is wide open between the other three teams - Poland and Ukraine are a similar standard to Romania and Hungary, who we came up against in qualifying.

*GPT3*: Germany will be the favourites, but it is

wide open between the other three teams. [44.09%, in range]

REFEREE-CONTROL: Germany will be the favourites, but it is wide open between the other three teams. [44.09%, in range]

*Ghalandari et al. (2022)*: Germany will be the favourites but it is wide open between the other three teams - Poland and Ukraine. [54.84%, in range]

*Liu et al. (2022b)*: wide between the other three teams poland and are romania hungary we came [39.25%, 0.75% too short]

*Original*: It is the latest mass killing attributed to the extremists, activists said, with the violence coming during a 24-hour period.

*GPT3*: The latest mass killing attributed to the extremists occurred during a 24-hour period. [68.8%, 8.8% too long]

REFEREE-CONTROL: It is the latest mass killing, activists said, with the violence coming. [57.6%, in range]

*Ghalandari et al. (2022)*: It is the latest mass killing attributed to the extremists 24-hour. [53.6%, in range]

*Liu et al. (2022b)*: is the latest mass killing attributed to the extremists [44.0%, in range]

*Original*: The brutal honesty and reality of death speaks to the truth and tragedy of bad timing for these two characters."

*GPT3*: The brutal honesty and reality of death speaks to the truth and tragedy of bad timing for these two characters. [99.11%, 39.11% too long]

REFEREE-CONTROL: The brutal honesty and reality of death speaks to the truth. [53.57%, in range]

*Ghalandari et al. (2022)*: The brutal honesty and of death speaks to timing characters. [53.57%, in range]

*Liu et al. (2022b)*: the reality of death the truth and of bad timing characters [52.68%, in range]

*Original*: It's a nice touch of confectionerial twinning, and broadly chimes with the optional theme of this year's show, 'Man Made', which judging by the parrots and skyscapes and severed heads and moths has not been taken up with widespread enthusiasm.

*GPT3*: The show's optional theme is 'Man Made'. [16.46%, 23.54% too short]



REFEREE-CONTROL: It's a nice touch of confectionery and broadly chimes with the optional theme of this year's show, "Man Made", which has not been taken up. **[57.2%, in range]**

*Ghalandari et al. (2022)*: It's a nice touch of confectionery twinning and chimes with theme of this year's show 'Man Made parrots. **[43.62%, in range]**

*Liu et al. (2022b)*: it 's a nice touch of <lunkl> severed **[15.23%, 24.77% too short]**

*Original*: Vicis, a finalist for the 2016 GeekWire Awards in the Hardware/Gadget of the Year category, will use the fresh cash for production and commercialization of its ZERO1 helmet for the upcoming football season.

*GPT3*: Vicis, a finalist for the 2016 GeekWire Awards in the Hardware/Gadget of the Year category, will use the fresh cash for production and commercialization of its ZERO1 helmet. **[83.98%, 23.98% too long]**

REFEREE-CONTROL: Vicis, a finalist for the 2016 GeekWire Awards, will use the fresh cash for production and commercialization. **[52.91%, in range]**

*Ghalandari et al. (2022)*: Vicis, a finalist for 2016 GeekWire Awards/Gadget will use fresh cash for its ZERO1 helmet. **[44.17%, in range]**

*Liu et al. (2022b)*: a finalist for the of the year helmet the season **[23.3%, 16.7% too short]**

*Original*: The cloud-based system has to wait for a connection to download updates to apps and content, but the content also has to be stored locally for uninterrupted playback.

*GPT3*: The cloud-based system has to wait for a connection to download updates to apps and content, but the content also has to be stored locally for uninterrupted playback. **[100.0%, 40.0% too long]**

REFEREE-CONTROL: The cloud-based system has to wait for a connection to download updates to apps and content. **[55.42%, in range]**

*Ghalandari et al. (2022)*: The cloud-based system has to wait for a connection to download updates to apps. **[48.19%, in range]**

*Liu et al. (2022b)*: the <lunkl> system has to wait for a download content stored locally **[40.96%, in range]**

*Original*: It signals toughness and finality, yet its

meaning is elusive.

*GPT3*: The article does not provide a clear explanation of what the phrase "yet its meaning is elusive" means. **[166.13%, 106.13% too long]**

REFEREE-CONTROL: It signals toughness and finality. **[54.84%, in range]**

*Ghalandari et al. (2022)*: It signals toughness and finality **[53.23%, in range]**

*Liu et al. (2022b)*: it toughness its meaning is elusive **[56.45%, in range]**

## E.2 Random sample of generations used in the 20-40% human evaluation

*Original*: Viggo left South America aged 11, when his parents divorced and he moved to upstate New York with his mother and siblings.

*GPT3*: Viggo left South America when he was 11. **[32.79%, in range]**

REFEREE-CONTROL: Viggo left South America aged 11, when his parents. **[41.8%, 1.8% too long]**

*Ghalandari et al. (2022)*: Viggo left South America aged 11 parents divorced. **[40.98%, 0.98% too long]**

*Liu et al. (2022b)*: parents moved to upstate new york **[27.05%, in range]**

*Original*: And Yahoo appears to have had another money-losing fumble on its hands with the free livestream of the NFL's Oct. 25 Bills-Jaguars game from London.

*GPT3*: Yahoo lost money streaming the NFL game from London. **[35.14%, in range]**

REFEREE-CONTROL: Yahoo appears to have had another money-losing fumble. **[36.49%, in range]**

*Ghalandari et al. (2022)*: And Yahoo had another money-losing fumbles Bills-Jaguars **[37.84%, in range]**

*Liu et al. (2022b)*: yahoo to have hands the of the nfl **[22.97%, in range]**

*Original*: Suárez almost levelled with his latest sweeping free-kick from distance but the clearer chances came to Hull against a ponderous Liverpool defence.

*GPT3*: Hull scored against Liverpool. **[20.41%, in range]**

REFEREE-CONTROL: Suarez almost levelled with his latest free-kick. **[33.33%, in range]**

*Ghalandari et al. (2022)*: Suárez almost levelled

with free-kick Hull Liverpool [35.37%, in range]  
*Liu et al. (2022b)*: distance chances came to hull  
liverpool [26.53%, in range]

*Original*: Specifically, the company agreed to limit pre-payments, to provide accurate estimates of charges, and to disclose details of financing agreements.

*GPT3*: The company agreed to limit pre-payments, to provide accurate estimates of charges, and to disclose details of financing agreements. [90.41%, 50.41% too long]

REFEREE-CONTROL: The company agreed to limit pre-payments and to provide accurate estimates. [51.37%, 11.37% too long]

*Ghalandari et al. (2022)*: Specifically company agreed to limit pre-payments. [34.25%, in range]

*Liu et al. (2022b)*: company to provide accurate disclose financing [31.51%, in range]

*Original*: "There was a family in the car, they got out before it hit.

*GPT3*: There was a family in the car, they got out before it hit. [98.31%, 58.31% too long]

REFEREE-CONTROL: "There was a family." [35.59%, in range]

*Ghalandari et al. (2022)*: "There was family car [35.59%, in range]

*Liu et al. (2022b)*: there was a it [23.73%, in range]

*Original*: As for the coming revolution, I can't say for certain when that will happen.

*GPT3*: I don't know when the revolution will happen. [59.21%, 19.21% too long]

REFEREE-CONTROL: I can't say for sure. [27.63%, in range]

*Ghalandari et al. (2022)*: As for the coming revolution [36.84%, in range]

*Liu et al. (2022b)*: the revolution i can certain [36.84%, in range]

*Original*: In one New Jersey county, the Red Cross was AWOL, says Lt. Matthew Tiedemann of the Bergen County Office of Emergency Management.

*GPT3*: The Red Cross was not present in one New Jersey county. [42.64%, 2.64% too long]

REFEREE-CONTROL: In one New Jersey county, the Red Cross was AWOL. [37.98%, in range]

*Ghalandari et al. (2022)*: In New Jersey Red Cross was AWOL [24.81%, in range]

*Liu et al. (2022b)*: new jersey county red cross awol lt. [27.91%, in range]

*Original*: Furthermore, though some might argue that the purpose of food is merely to remind us of our own mortality, I would argue that the best cuisine can be an art form.

*GPT3*: I think that food can be an art form. [22.84%, in range]

REFEREE-CONTROL: Furthermore, I would argue that the best cuisine can be an art form. [41.98%, 1.98% too long]

*Ghalandari et al. (2022)*: Furthermore of food mortality the best cuisine art form. [34.57%, in range]

*Liu et al. (2022b)*: our mortality the best cuisine can be an art [27.16%, in range]

*Original*: Without proper conservation facilities nearby, the artifacts Batchvarov and his colleagues discovered had to be reburied underwater.

*GPT3*: The artifacts Batchvarov and his colleagues discovered had to be reburied underwater. [64.39%, 24.39% too long]

REFEREE-CONTROL: The artifacts had to be reburied. [25.0%, in range]

*Ghalandari et al. (2022)*: Without artifacts Batchvarov reburied underwater [36.36%, in range]

*Liu et al. (2022b)*: artifacts to be reburied underwater [26.52%, in range]

*Original*: Often people are desperate or ridiculous rather than heroic, so I do what needs to be done."

*GPT3*: The protagonist does what needs to be done. [46.74%, 6.74% too long]

REFEREE-CONTROL: I do what needs to be done. [29.35%, in range]

*Ghalandari et al. (2022)*: Often people are desperate heroic. [36.96%, in range]

*Liu et al. (2022b)*: do what to be done [19.57%, 0.43% too short]

*Original*: Liam's feet make a guest appearance in the picture, as does his Great Dane Watson who he bought with ex-girlfriend Sophia Smith can be seen by his settee.

*GPT3*: There are pictures of Liam and his dog Watson in the article. [39.61%, in range]

REFEREE-CONTROL: Liam's feet make a guest appearance in the picture. [33.12%, in range]

*Ghalandari et al. (2022)*: Liam's feet make guest appearance in picture Dane Watson girlfriend. [43.51%, 3.51% too long]

*Liu et al. (2022b)*: dane watson who bought ex-girlfriend sophia smith can be [36.36%, in range]

*Original*: I like this little gallery, the Small Weston Room, where the paintings go right up to the ceiling.

*GPT3*: I like this little gallery. [27.55%, in range]

*REFEREE-CONTROL*: I like this little gallery. [27.55%, in range]

*Ghalandari et al. (2022)*: I like gallery Small Weston Room paintings [42.86%, 2.86% too long]

*Liu et al. (2022b)*: go right to the ceiling [23.47%, in range]

*Original*: After a stint as a lobbyist, during which he helped establish two think tanks affiliated with James Dobson's hugely influential Focus on the Family, Parker became a confidant of Roy Moore, then a county judge.

*GPT3*: Parker became a confidant of Roy Moore and helped him win election as a county judge. [40.67%, 0.67% too long]

*REFEREE-CONTROL*: Parker became a confidant of Roy Moore, then a county judge, during a stint. [36.36%, in range]

*Ghalandari et al. (2022)*: After as lobbyist Dobson, Parker became a confidant of Roy Moore. [31.1%, in range]

*Liu et al. (2022b)*: a a lobbyist dobson focus the family parker became a [24.88%, in range]

*Original*: Germany will be the favourites but it is wide open between the other three teams - Poland and Ukraine are a similar standard to Romania and Hungary, who we came up against in qualifying.

*GPT3*: Germany is the favourite to win the World Cup, but it's still wide open. [38.71%, in range]

*REFEREE-CONTROL*: Germany is the favourite, but it is wide open between the other three teams. [40.86%, 0.86% too long]

*Ghalandari et al. (2022)*: Germany will be the favourites but between teams Poland Ukraine. [34.41%, in range]

*Liu et al. (2022b)*: teams poland and ukraine standard romania hungary we came qualifying [36.56%, in range]

*Original*: It is the latest mass killing attributed to the extremists, activists said, with the violence coming during a 24-hour period.

*GPT3*: Extremists have killed more people in 24 hours than ever before. [51.2%, 11.2% too long]

*REFEREE-CONTROL*: It is the latest mass killing, activists said. [36.8%, in range]

*Ghalandari et al. (2022)*: It is latest mass killing attributed extremists [37.6%, in range]

*Liu et al. (2022b)*: latest mass killing attributed to extremists [35.2%, in range]

*Original*: The brutal honesty and reality of death speaks to the truth and tragedy of bad timing for these two characters.'

*GPT3*: Death speaks to the truth and tragedy of bad timing in the characters' lives. [68.75%, 28.75% too long]

*REFEREE-CONTROL*: The brutal honesty and reality of death. [35.71%, in range]

*Ghalandari et al. (2022)*: The brutal honesty death speaks characters. [38.39%, in range]

*Liu et al. (2022b)*: the reality the truth timing characters [34.82%, in range]

*Original*: It's a nice touch of confectionerial twinning, and broadly chimes with the optional theme of this year's show, 'Man Made', which judging by the parrots and skyscapes and severed heads and moths has not been taken up with widespread enthusiasm.

*GPT3*: The exhibit has a theme. [9.88%, 10.12% too short]

*REFEREE-CONTROL*: It's a nice touch of confectionery, and it's also a nice theme of this year's show. [34.16%, in range]

*Ghalandari et al. (2022)*: It's a nice touch of confectionerial twinning show 'Man Made. [25.1%, in range]

*Liu et al. (2022b)*: it 's a touch of the [8.23%, 11.77% too short]

*Original*: Vicis, a finalist for the 2016 GeekWire Awards in the Hardware/Gadget of the Year category, will use the fresh cash for production and commercialization of its ZERO1 helmet for the upcoming football season.

*GPT3*: Vicis will use the fresh cash to produce and commercialize its ZERO1 helmet. [36.89%, in range]

REFEREE-CONTROL: Vicis will use the fresh cash for production and commercialization. [**32.52%**, **in range**]

*Ghalandari et al. (2022)*: Vicis finalist GeekWire Awards/Gadget will use cash its ZERO1 helmet [**33.01%**, **in range**]

*Liu et al. (2022b)*: a finalist the of the year category helmet the season [**25.73%**, **in range**]

*Original*: The cloud-based system has to wait for a connection to download updates to apps and content, but the content also has to be stored locally for uninterrupted playback.

*GPT3*: Cloud-based systems have to wait for updates, but the content can be stored locally for uninterrupted playback. [**66.87%**, **26.87% too long**]

REFEREE-CONTROL: The cloud-based system has to wait for a connection. [**31.33%**, **in range**]

*Ghalandari et al. (2022)*: The cloud-based system has to wait for download updates. [**33.73%**, **in range**]

*Liu et al. (2022b)*: <lunk> has wait for a download content locally [**28.31%**, **in range**]

*Original*: It signals toughness and finality, yet its meaning is elusive.

*GPT3*: The sign means something but we don't know what it is. [**87.1%**, **47.1% too long**]

REFEREE-CONTROL: It's a mystery. [**24.19%**, **in range**]

*Ghalandari et al. (2022)*: It signals toughness finality [**46.77%**, **6.77% too long**]

*Liu et al. (2022b)*: signals toughness elusive [**40.32%**, **0.32% too long**]