

---

# Appendices for A Simple Contrastive Learning Objective for Alleviating Neural Text Degeneration

---

Anonymous Author(s)

Affiliation

Address

email

## 1 A Ethical considerations

2 In this work, we used publicly available English data to train/validate/test models. As far as we  
3 know, the curators of these datasets have taken ethical issues into consideration when creating the  
4 datasets. We manually checked some generated texts of the language models trained by CT and  
5 did not observe any noticeable traces of concern, such as offensive and malevolent language. We  
6 share our source code and trained model weights to support its correct use. To make sure the human  
7 workers involved in the data labeling efforts, as part of the human evaluation for this study, are fairly  
8 paid, we applied the minimum hourly rate of 10.48 euros, which converts to 11 dollars per hour.  
9 However, we warn that generative language models should always be used with caution since the  
10 generated texts are usually novel and unexpected wordings may appear when trained on improper  
11 data. Especially, generative models can be used maliciously, e.g., to generate fake news articles.

## 12 B Using CT in your work

---

**Algorithm 1** Calculate contrastive token loss

---

**Input:** Labels  $X = (x_1, x_2, \dots, x_{|X|})$ , time  $t$ , negative window size  $M$ , logits  $Z_t$  of time  $t$

**Output:** Contrastive token loss  $\mathcal{L}_{CT}^t$

1:  $S_N^t \leftarrow \text{SampleNegatives}(X, M, t)$  # according to Eq. (7)  
2:  $z_{x_t} \leftarrow \text{GatherLogits}(Z_t, x_t)$  # positive logits  
3:  $z_{S_N^t} \leftarrow \text{GatherLogits}(Z_t, S_N^t)$  # negative logits  
4:  $\mathcal{L}_{CT}^t \leftarrow \log \left( 1 + \sum_{x_i^- \in S_N^t} \exp(z_{x_i^-} - z_{x_t}) \right)$  # Eq. (5)  
5: **return**  $\mathcal{L}_{CT}^t$

---

13 We summarize the steps for calculating  $\mathcal{L}_{CT}^t$  in Algorithm 1. You can use our CT objective when *pre-*  
14 *training* or *finetuning* your autoregressive language models, which takes only several lines of Python  
15 code, around where you calculate PyTorch’s CrossEntropyLoss. Simply use `pip install ct-loss` to  
16 install the required packages. Then you can use CT as follows:

```
17 1 import torch
18 2
19 3 # Suppose we already have the model output logits and labels (sequences
20 4 # of token indices). For example when the batch size is 10, sequence
21 5 # length is 50 and vocabulary size is 1000:
22 6 logits = torch.rand(10, 50, 1000)
23 7 labels = torch.randint(0, 999, (10, 50))
24 8
25 9 # This is how you normally use cross-entropy for a language model:
```

```

26|0 from torch.nn import CrossEntropyLoss
27|1 ce_criterion = CrossEntropyLoss()
28|2 ce_loss = ce_criterion(logits.view(-1, 1000), labels.view(-1))
29|3
30|4 # This is how you can use our contrastive token loss:
31|5 from ct.ct_loss import ContrastiveTokenLoss
32|6 ct_criterion = ContrastiveTokenLoss(pad_id=999) # we need pad tokens
33|7     for masking out tokens in a sequence that should not be used as
34|8     negative tokens
35|9 ct_loss = ct_criterion(logits, labels)
36|0
37|1 # In our paper, we use CE and CT together
38|2 loss = ce_loss + ct_loss

```

## 39 C Noise-contrastive estimation for autoregressive language models

40 We adapted NCE [2] to token-level:

$$\mathcal{L}_{NCE}^t = -\log \sigma(h_t^T W_{x_t}) - \frac{1}{|S_N^t|} \sum_{x_t^- \in S_N^t} \log \sigma(-h_t^T W_{x_t^-}), \quad (8)$$

41 where  $\sigma(\cdot)$  is the *sigmoid* function.

## 42 D Gradient functions

43 To see how loss functions influence the logits during training, we compare the gradient of each loss  
44 function. Writing  $z_{x_t} = h_t^T W_{x_t}$  for the logit of token  $x_t$ , the gradient function is calculated by  
45  $\partial \mathcal{L}_* / \partial z_*$ , where  $\mathcal{L}_* \in \{L_{CE}, L_{UL}, L_{CT}\}$ , and  $z_* \in \{z_{x_t}, z_{\hat{x}_t}, z_{x_t^-}\}$ . For clarity, we further denote  
46  $p(*|x_{<t})$  as  $p_*$ .

47 • Gradient functions of cross-entropy, w.r.t. label tokens  $x_t$ :

$$\begin{aligned} \frac{\partial \mathcal{L}_{CE}}{\partial z_{x_t}} &= -\frac{\sum_{\hat{x}_t \in V, \hat{x}_t \neq x_t} \exp(z_{\hat{x}_t} - z_{x_t})}{1 + \sum_{\hat{x}_t \in V, \hat{x}_t \neq x_t} \exp(z_{\hat{x}_t} - z_{x_t})} \\ &= -\frac{\sum_{\hat{x}_t \in V, \hat{x}_t \neq x_t} \exp(z_{\hat{x}_t})}{\exp(z_{x_t}) + \sum_{\hat{x}_t \in V, \hat{x}_t \neq x_t} \exp(z_{\hat{x}_t})} \\ &= -\sum_{\hat{x}_t \in V, \hat{x}_t \neq x_t} p_{\hat{x}_t} \\ &= p_{x_t} - 1 \\ &\leq 0, \end{aligned} \quad (9)$$

48 and non-label tokens  $\hat{x}_t$  (including negative tokens and irrelevant tokens):

$$\begin{aligned} \frac{\partial \mathcal{L}_{CE}}{\partial z_{\hat{x}_t}} &= \frac{\exp(z_{\hat{x}_t} - z_{x_t})}{1 + \sum_{\hat{x}_t \in V, \hat{x}_t \neq x_t} \exp(z_{\hat{x}_t} - z_{x_t})} \\ &= \frac{\exp(z_{\hat{x}_t})}{\exp(z_{x_t}) + \sum_{\hat{x}_t \in V, \hat{x}_t \neq x_t} \exp(z_{\hat{x}_t})} \\ &= p_{\hat{x}_t} \\ &\geq 0. \end{aligned} \quad (10)$$

- 49 • Gradient functions of unlikelihood training w.r.t. negative tokens  $x_t^-$ :

$$\begin{aligned}
\frac{\partial \mathcal{L}_{UL}}{\partial z_{x_t^-}} &= - \sum_{x_t^- \in C^t} \frac{\partial \log(1 - p_{x_t^-})}{\partial p_{x_t^-}} \frac{\partial p_{x_t^-}}{\partial z_{x_t^-}} \\
&= \sum_{x_t^- \in C^t} \frac{1}{1 - p_{x_t^-}} \frac{\partial p_{x_t^-}}{\partial z_{x_t^-}} \\
&= p_{x_t^-} - \sum_{x_t'^- \in C^t, x_t'^- \neq x_t^-} \frac{p_{x_t^-} p_{x_t'^-}}{1 - p_{x_t'^-}} \\
&= p_{x_t^-} \left( 1 - \sum_{x_t'^- \in C^t, x_t'^- \neq x_t^-} \frac{p_{x_t'^-}}{1 - p_{x_t'^-}} \right) \\
&\in (-\infty, p_{x_t^-}],
\end{aligned} \tag{11}$$

- 50 and other tokens  $\hat{x}_t$  (including label tokens and irrelevant tokens):

$$\begin{aligned}
\frac{\partial \mathcal{L}_{UL}}{\partial z_{\hat{x}_t}} &= - \sum_{x_t^- \in C^t} \frac{\partial \log(1 - p_{x_t^-})}{\partial p_{x_t^-}} \frac{\partial p_{x_t^-}}{\partial z_{\hat{x}_t}} \\
&= \sum_{x_t^- \in C^t} \frac{1}{1 - p_{x_t^-}} (-p_{x_t} p_{x_t^-}) \\
&= \sum_{x_t^- \in C^t} \frac{p_{x_t} p_{x_t^-}}{p_{x_t^-} - 1} \\
&\leq 0.
\end{aligned} \tag{12}$$

- 51 • Gradient functions of CT w.r.t. positive tokens  $x_t$ :

$$\begin{aligned}
\frac{\partial \mathcal{L}_{CT}}{\partial z_{x_t}} &= - \frac{\sum_{x_t^- \in S_N^t} \exp(z_{x_t^-} - z_{x_t})}{1 + \sum_{x_t^- \in S_N^t} \exp(z_{x_t^-} - z_{x_t})} \\
&= - \frac{\sum_{x_t^- \in S_N^t} p_{x_t^-} / p_{x_t}}{1 + \sum_{x_t^- \in S_N^t} p_{x_t^-} / p_{x_t}} \\
&\leq 0,
\end{aligned} \tag{13}$$

- 52 and negative tokens  $x_t^-$ :

$$\begin{aligned}
\frac{\partial \mathcal{L}_{CT}}{\partial z_{x_t^-}} &= \frac{\exp(z_{x_t^-} - z_{x_t})}{1 + \sum_{x_t'^- \in S_N^t} \exp(z_{x_t'^-} - z_{x_t})} \\
&= \frac{p_{x_t^-} / p_{x_t}}{1 + \sum_{x_t'^- \in S_N^t} p_{x_t'^-} / p_{x_t}} \\
&\geq 0.
\end{aligned} \tag{14}$$

- 53 Because all terms in Eq. (5) are independent with irrelevant tokens  $\hat{x}_t$ :

$$\frac{\partial \mathcal{L}_{CT}}{\partial z_{\hat{x}_t}} = 0. \tag{15}$$

- 54 • NCE with respect to label tokens  $x_t$ :

$$\begin{aligned}
\frac{\partial \mathcal{L}_{NCE}}{\partial z_{x_t}} &= -\sigma(z_{x_t})(1 - \sigma(z_{x_t})) \\
&\leq 0,
\end{aligned} \tag{16}$$

55 and negative tokens  $x_t^-$ :

$$\frac{\partial \mathcal{L}_{NCE}}{\partial z_{x_t^-}} = \sigma(-z_{x_t^-})(1 - \sigma(-z_{x_t^-})) \geq 0. \quad (17)$$

56 Same as CT, all terms in Eq. (8) are independent with irrelevant tokens  $\hat{x}_t$ :

$$\frac{\partial \mathcal{L}_{NCE}}{\partial z_{\hat{x}_t}} = 0. \quad (18)$$

## 57 E Required software and hardware resources

58 For the CE and decoding baselines, we use GPT-2 [4] implemented and pretrained using the CE  
 59 objective by Hugging Face [12]. For fair comparisons, we implement our CT loss and all learning-  
 60 based baselines and use them to train GPT-2. Specifically, for unlikelihood training, we implemented  
 61 both the token-level (UL-T) and the sequence-level (UL-S) variants, according to the official source  
 62 code [10]. We also implemented SimCTG according to the official code [8]. Similar to CT, we  
 63 adapted NCE to the token-level (detailed in Appendix In our experiments, NCE is also used together  
 64 with CE as was done for CT in Eq. (6).

65 Our implementation is based on Hugging Face Transformers (Apache-2.0 license) [12], PyTorch  
 66 Lightning (Apache-2.0 license) [11], and Hydra (MIT license) [13]. Our source code is directly  
 67 based on Lightning Transformers (Apache-2.0 license) [9], thus inheriting the license. All our ex-  
 68 periments are conducted on a single TITAN Xp GPU and use less than 20GB of CPU memory.

## 69 F Additional results and analysis for the language modeling task

### 70 F.1 Additional results

71 Figure 4 reveals that the heat maps for NCE, UL-T and SimCTG are similar to that of CE in Figure  
 72 3. More specifically, they all contain excessive stripes, although less so with NCE due to its lower  
 73 repetition rates. Besides, they are also darker at the lower-right half of the diagonal cells, especially  
 for NCE and SimCTG.




Figure 4: Heat maps for the generation probability of NCE, UL-T and SimCTG on the Wikitext-103 test set.

74

75 Table 5 showcases the *ungrammatical token repetition* problem of UL-TS when trained using a  
 76 larger learning rate of  $1e-5$ , while it is not a problem with CT trained using a learning rate of  $1e-4$ .  
 77 In Table 6, we show more examples of comparing the generated texts of CT with those by other  
 78 approaches.

### 79 F.2 Breakdown analysis

80 Beyond the overall performance analysis given above, we also provide a breakdown analysis for CT.



Table 6: More examples of CT compared to other approaches.

|           |  | rep-1 |
|-----------|--|-------|
| Prefix    | deep load. During the ships' modernization during the 1930s, their forward superstructures were enlarged with multiple platforms added to their tripod foremasts. Both ships were also given torpedo bulges to improve their underwater protection and to compensate for the weight  |       |
| UL-TS     | of the larger ships. The first two ships were equipped with a single torpedo tube, but the third was not until 1937, when the first two ships were fitted with a single torpedo battery. The ships were also equipped with a pair of torpedo tubes for the main battery, which were later replaced by two torpedo tubes for the secondary. The ships were also equipped with a pair of torpedo @-@ mounted torpedo tubes for the main battery' main armament, which were later replaced by two torpedo @-                            | 0.58  |
| SimCTG-CS | of the ship's armor. = = Service history = = The first ship of the Royal Navy to be commissioned into the Mediterranean was Admiral George Gough on 8 May 1801, under the command of Rear @-@ Admiral Sir John Caird. Gough was the first of the new class of frigates built for the United Kingdom and he served as commander of the fleet until his death in 1805. His flagship, HMS Illustrious, was laid down at Portsmouth in March 18  | 0.3   |
| Top-k     | of their forward bulkheads. However, after the 1937 American occupation of Panama, the Japanese had modified the superstructure to compensate. = = = Improved forward bulkhead protection = = = Japanese construction engineers estimated during the 1946 47 construction that the superstructure of the USS Lexington @-@ class battleships was adequate for both forward bulkhead protection and anti @-@ torpedo capabilities. The armor on the bow was 5 @. @ 5 inches thick and the rear armor had a                            | 0.4   |
| CT        | loss caused by heavy seas. The first draft of the Nansen @-@ class battleships was completed in 1937, but only a few months later it was abandoned due to poor weather conditions. = = Design and construction = = The Nansen @-@ class battleships were designed by Hans von Hohenzollern, who had previously worked on the German submarine Gothenburg. The ship's main armament consisted of four 45 mm ( 2 @. @ 9 in )   | 0.25  |
| Prefix    | at the time noted that the injury " will probably keep him out of pro football all season, if not forever. " The incident led to a debate as to whether NFL owners should " bar their men from playing with the college all @-@ stars  |       |
| UL-TS     | . " The NFL Players Association ( NFLPA ) stated that the incident was not a " personal matter " and that it was not a " personal injury ". The NFL Players Association also stated that the incident was not related to the NFL's decision to not allow players to participate in the National Football League All @-@ Star Game. = = = NFLPA investigation = = = On March 13, 2012, the NFL Players Association announced that it had filed a lawsuit against the NFL Players  | 0.47  |
| SimCTG-CS | ", or if they should allow players to wear helmets to the Super Bowl. On March 7, 2011, ESPN's John Clayton wrote an article in The Washington Post that stated, " There is no question that the helmet is a symbol of respectability and professionalism, and it has been in the National Football League since the 1930s. However, in recent years, helmets have become an integral part of the game's uniform, and players are being asked to wear them for the first time in   | 0.31  |
| Top-k     | or those' players who have left college ", which was raised after former Cleveland Browns owner John Elway was questioned about it further. Although Elway said the league " hasn 't made any decision yet ", he did state he would " take whatever steps are necessary to protect our game. " Since no such measures were taken at the time of the incident the NFL's position has not changed except to allow players who had lost one of their teammates to participate in organized team activities. Elway                       | 0.23  |
| CT        | or simply allow them to retire at the end of the season. " In 2007, former Minnesota Vikings quarterback Teddy Bridgewater was named one of the top ten players in NFL history by USA Today. = = Early life = = Born in St. Louis, Missouri, Cameron Reed was raised by his mother and grandmother. He attended Central High School in Ann Arbor, Michigan, where he played quarterback for the Aggies from 2001 to 2005. After graduating from Central High School, Reed moved to Los Angeles                                       | 0.28  |
| Prefix    | Doniel decided to abandon the position that night. When it got dark the survivors would split into small groups and try to get back to friendly lines. That evening after dark the North Koreans launched another weak attack against the position. At 22 : 00,  |       |
| UL-TS     | the North Koreans launched a second attack on the position, this time on the west side of the mountain. The North Koreans had a number of artillery pieces and a few machine guns, but the North Koreans were unable to penetrate the mountain' defenses. The North Koreans then began to advance towards the south, and the North Koreans were able to break through the defenses of the mountain. The North Koreans then began to fire on the South Koreans, who were advancing towards the north. The North Koreans then began to | 0.54  |
| SimCTG-CS | a U @-@ boat was sunk by an American submarine and two others were damaged, one of which was badly damaged in the crash. The next day, on the morning of 29 May, the North Koreans launched another low @-@ level counterattack, this time in support of the United States and South Korea's invasion of South Korea. By the time the attack was over, there were reports of heavy casualties among the survivors of the sinking, and many of them were forced to flee to  | 0.34  |
| Top-k     | a group of 13 North Korean artillerymen was hit. At 23 : 55, an attack was launched on the southern flank of the column. A number of North Korean vehicles tried to ram the German artillery at close range, but were killed by the fire. All the tanks in that column were eliminated by the German sides. Only the small tanks and two armoured personnel carriers were damaged. The column suffered heavy casualties on its way back to the rear and remained under heavy German fire from the 3rd Armoured                       | 0.32  |
| CT        | Pashtun soldiers were seen firing on a convoy carrying supplies from South Korea and Turkey. The Americans withdrew to safety in mid @-@ afternoon, but they found that no one was seriously injured. = = Battle of Chongju Island = = On 9 August 1945, U.S. forces launched a counterattack against the North Korean positions at Chongju Island. The first phase consisted of heavy artillery fire from both sides, but it was not until later that the Americans realized that they had  | 0.23  |

## Excerpt

... be a monophyletic group, and sister to the clade containing Allagoptera, Polyandrococos, Parajubaea, Butia and Jubaea. Disagreement exists as to whether Attalea should be considered

### Continuation 1

a single genus, or a group of related genera. In their 1996 Field Guide to the Palms of the Americas, Andrew Henderson, Gloria Galeano and Rodrigo Bernal combined all the species in the subtribe Attaleinae ( as it was then defined ) into a single genus, Attalea. In his 1999 Taxonomic Treatment of Palm Subtribe Attaleinae, American botanist Sidney F. Glassman divided the group into five genera — a more narrowly defined ...

### Continuation 2

a separate species from its closest relatives. The current definition of " Nativist " refers to those who believe that plants are inherently beautiful and thus deserving of protection from predators, whereas others consider them merely decorative objects. In contrast, some authors have argued that Attalea's ability to reproduce naturally is due to its unique genetic makeup. = = Description = = The fruit bodies are cylindrical with a width of about 2 @ .@ 5 cm ( 1 @ .@ 8 in ...

Which continuation is **less repetitive**:

- Continuation 1    Continuation 2    Not sure

Which continuation is **more fluent**:

- Continuation 1    Continuation 2    Not sure

Which continuation is **more coherent**:

- Continuation 1    Continuation 2    Not sure

**In all**, which continuation do you think is better:

- Continuation 1    Continuation 2    Not sure

Please justify your answers:

Submit

Figure 7: Our MTurk question form design for the human evaluation on the language modeling task.

## 94 G Human evaluation design

95 Figure 7 is a screen shot of our design of question form. We instructed the crowd workers to first  
96 read the excerpt (prefix to LMs) and the generated continuations, and then to compare their quality  
97 from three aspects: repetitiveness, fluency and coherence. We allow the workers to choose “Not  
98 sure” when they cannot tell which continuation is better. Based on their answers, the workers were  
99 also asked to select the overall winner. For quality control, we also asked the workers to provide a  
100 justification message. Please see Figure 8 for the full instruction.

## 101 H Experimental setup for the dialogue task

102 The experimental setup for the dialogue task below follows largely that of the language modeling  
103 task in §5. Below we focus on the differences.

104 **Datasets.** We follow Roller et al. [6] to use a mixture of multiple high-quality datasets, including  
105 PersonaChat [14], Empathetic Dialogues [5], Wizard of Wikipedia [1], and BlendedSkillTalk [7].  
106 We add another benchmark dialogue dataset DailyDialog [3]. For each training example, we use up  
107 to 3 turns of dialogue history as the input context, and 1 follow-up turn as the target response.

108 **Training and Inference Details.** We use the *400M-distilled* version BlenderBot [6] implemented  
109 and pretrained using the CE objective by Hugging Face [12]. We truncate the maximum of sequence

## Select the better text continuation

We are researchers working on natural language generation. Our sincere thanks to you for helping out. In this HIT, you will see a human-written text excerpt from Wikipedia, and two continuations that may be generated by human or computer programs. These continuations should continue writing from the end of the excerpt. Your task is to compare which continuation fits better with the excerpt.

### Instructions

After reading the excerpt and continuations, you need to compare the quality of the continuations from three aspects: **repetitiveness**, **fluency** and **coherence**. The better continuation is the one that's less repetitive, more fluent and more coherent, and we provide one question for each aspect. We ask you to choose a winner for each of these aspects. When they look equally good/bad on one aspect, you can answer **Not sure** for the corresponding question. Sometimes, it's hard for one continuation to win all three aspects, then you need to decide which one wins more. If finally both continuations look equally good/bad, on all three aspects, you can also answer *Not sure* for the 4-th question (the overall quality). We treat all three aspects equally important.

You also need to write a specific justification for your answers, by providing proofs from the excerpt and/or continuations, and explain how they support your answers. Failure to do so will result in your answer being rejected.

### Examples

To help you better understand the three aspects, we provide some examples below.

The following sentence is **repetitive**, as highlighted:

The poem's themes are often divided into three main themes : the " dark ", " light @-@ hearted " and " light @-@ hearted ".

The following sentence is **not fluent** because usually you wouldn't take a ship to the hospital, neither will you break it up in there:

Two days later, the ship was attacked by a group of U @-@ boats and sank with no survivors. She was taken to a hospital and later broken up for scrap.

The following example is **incoherent**. In the HIT you may see incoherent information between the continuation and the excerpt, or within a continuation itself.

A few days later, two of the survivors are killed in the accident, one of whom is taken to a hospital where he is treated for burns on his face and hands. He later becomes a member of...

In contrast, here is a good example (at least we believe so, because we selected from real Wikipedia data):

**Excerpt:** ... = = Meteorological history = = The origins of the hurricane were from a tropical wave that possibly spawned a tropical depression on August 27, although there **Continuation:** was minimal data over the next few days as it tracked to the west @-@ northwest. On August 31, a nearby ship reported gale force winds, which indicated that a tropical storm had developed to the east @-@ northeast of the Lesser Antilles. Based on continuity, it is estimated the storm attained hurricane status later that day. Moving quickly to the west @-@ northwest, the storm passed north of the Lesser Antilles and Puerto Rico...

When checked using our criteria, the above continuation is non-repetitive, fluent, and coherent with the excerpt as well as with itself. Therefore, we can say this is a good continuation.

**Please note** that some of the continuations were generated by computer programs, and these programs are not very precise with times, relationships of celebrities, etc. But don't bother checking their factuality, just feel by yourself if they make sense or not. The excerpt may occasionally end at a sub-word. E.g., the excerpt may end with "lakes" and the continuation begins with "ide", together they form the word "lakeside". There may also be some formatting symbols, most commonly they are "=" and "@", etc. Thanks again for contributing to this HIT.

Figure 8: Our instructions to MTurk workers.

110 length to 128 tokens, and a training batch of 10 context-response pairs. We follow Roller et al. [6]  
111 to force BlenderBot to generate at least 20 tokens.

## 112 I Results on the open-domain dialogue task

113 The results on the open-domain dialogue task are reported in Table 7. Generations have a minimum  
114 length of 20 tokens. Similar to its performance on the language modeling task, CT again achieves  
115 the best repetition and diversity performance, and with a minor sacrifice in terms of *ppl* (1.44 points).

116 Figure 9 indicates that CT has substantially more cases with lower repetition rates than other ap-  
117 proaches. Due to the fact that dialogue responses are usually short (~20 tokens), the *rep-4* rates  
118 of each method are not far apart, although CT marginally wins.

119 Regarding the selection of the sequence length for CT and the window size for selecting negative  
120 tokens, we made similar observations on the dialogue task as those on the language modeling task,  
121 as can be seen from Figure 10 and 11.

122 Table 8 shows some side-by-side comparisons of the responses generated by UL-TS and CT. One can  
123 observe that the dialogue responses generated by CT are usually less repetitive and more coherent  
124 with the on-going topics.

|                       | ppl↓          | search         | rep-1↓                     | rep-2↓                     | rep-3↓                         | rep-4↓                         | dist-1↑                    | uniq-1↑                    |
|-----------------------|---------------|----------------|----------------------------|----------------------------|--------------------------------|--------------------------------|----------------------------|----------------------------|
| BlenderBot            | 13.26         | greedy<br>beam | 25.77<br>13.34             | 12.17<br>3.56              | 8.23<br>2.01                   | 6.62<br>1.38                   | 0.56<br>0.62               | 5955<br>6144               |
| <i>decoding-based</i> | 3-gram ban    | greedy<br>beam | 20.30<br><b>11.13</b>      | 4.76<br><b>1.16</b>        | <i>0.00</i> ‡<br><i>0.00</i> ‡ | <i>0.00</i> ‡<br><i>0.00</i> ‡ | 0.57<br><b>0.62</b>        | 6031<br><b>6166</b>        |
|                       | Top- <i>k</i> | greedy<br>beam | <b>11.52</b><br>13.43      | <b>1.50</b><br>3.23        | <b>0.43</b><br><b>1.66</b>     | <b>0.23</b><br><b>1.05</b>     | <b>0.64</b><br>0.61        | <b>7043</b><br>6155        |
|                       | Nucleus       | greedy<br>beam | 13.04<br>13.61             | 2.17<br>3.35               | 0.81<br>1.76                   | 0.52<br>1.15                   | 0.62<br>0.61               | 6800<br>6138               |
|                       | SimCTG        | greedy<br>beam | 24.02<br>12.85             | 10.63<br>2.98              | 7.27<br>1.61                   | 6.15<br>1.10                   | 0.58<br>0.63               | 6171<br>6313               |
| <i>learning-based</i> | NCE           | greedy<br>beam | 14.40<br>9.53              | 2.50<br>1.20               | 0.88<br>0.42                   | 0.50<br>0.21                   | 0.59<br>0.62               | 6132<br>6122               |
|                       | UL-T          | greedy<br>beam | 21.02<br>10.64             | 8.80<br>2.02               | 6.23<br>0.93                   | 5.35<br>0.55                   | 0.57<br>0.63               | 6074<br>6204               |
|                       | UL-TS         | greedy<br>beam | 15.58<br>9.95              | 2.56<br>1.41               | 0.70<br>0.59                   | 0.28<br>0.29                   | 0.59<br>0.63               | 6209<br>6252               |
|                       | CT            | greedy<br>beam | <b>9.19</b><br><b>6.89</b> | <b>0.69</b><br><b>0.69</b> | <b>0.14</b><br><b>0.27</b>     | <b>0.05</b><br><b>0.12</b>     | <b>0.60</b><br><b>0.64</b> | <b>6404</b><br><b>6408</b> |
|                       | Human         | –              | –                          | 8.33                       | 0.83                           | 0.19                           | 0.06                       | 0.91                       |

Table 7: Results on the open-domain dialogue task. ‡ Does not count as the best.




Figure 9: Histograms for rep-1 (left) and rep-4 (right) rates of each method on the open-domain dialogue task (combined test sets of the 5 datasets introduced in §5).

## References

- 125
- 126 [1] Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston.  
127 2019. Wizard of wikipedia: Knowledge-powered conversational agents. In *7th International*  
128 *Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.  
129 OpenReview.net.
- 130 [2] Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estima-  
131 tion principle for unnormalized statistical models. In *Proceedings of the thirteenth interna-*  
132 *tional conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and  
133 Conference Proceedings.
- 134 [3] Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog:  
135 A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International*  
136 *Joint Conference on Natural Language Processing, IJCNLP 2017, Taipei, Taiwan, November*




Figure 10: Influence of sequence length for CT loss on the open-domain dialogue task.




Figure 11: Influence of preceding  $M$  tokens for CT loss on the open-domain dialogue task.

|         |   | rep-1 |
|---------|---|-------|
| Context | A: Oh really. It should be very nice to go there, listening to some music<br>B: Indeed. Have you ever heard of Pet sounds? Did you listen to any of their music?<br>A: I really haven't, but if is rock then I should like it, I always listen to some rock when I hike |       |
| UL-TS   | I love rock music. I also love folk music. I think I like the folk sound better.  | 0.41  |
| CT      | I love rock music. What else do you like to do? I am a big fan of the outdoors myself.  | 0.12  |
| Context | A: I am a kennel cleaner at a local animal shelter. And you?<br>B: I am a doctor in the labor and deliver unit.<br>A: Sounds very important. You must be a people person.   |       |
| UL-TS   | Yes I am. I also like to read. Do you? Do you like to read?   | 0.42  |
| CT      | Yes, I am. What do you do for fun? I like to hike.  | 0.16  |
| Context | A: Sounds like a nice way to make a little extra money<br>B: That is why I like getting up early so I can have more time for my crafts.<br>A: Well I am getting ready to turn in. Going to read some non fiction and fall asleep  |       |
| UL-TS   | I like to read too. I also like to play video games. What about you? What do you like to do?  | 0.40  |
| CT      | I like to read too. What are you going to read? I am currently reading the hunger games trilogy.  | 0.15  |
| Context | A: Hmm. Very interesting. Were you homeschooled? I was. You sing?<br>B: I do not sing and I was not homeschooled.<br>B: You ever been to california? What is ohio known for? You have teenagers?  |       |
| UL-TS   | I have never been to california. I am not a teenager. I have never been to the west coast   | 0.33  |
| CT      | No, I have not been to california. And no, I do not have teenagers.   | 0.25  |
| Context | A: Hi there. I am a hair dresser. Want to chat?   |       |
| UL-TS   | Sure. I am a hair stylist too. What do you do? I cut my hair short.   | 0.23  |
| CT      | Sure. What do you do for a living? I am a lawyer. What do you do?   | 0.43  |
| Context | Hello. How are you? I am in sale. What about you?   |       |
| UL-TS   | I am good. Just got back from the gym. What are you selling? I am a fitness instructor.   | 0.17  |
| CT      | I am good. What do you sell? I am a lawyer. What do you do?   | 0.42  |
| Context | A: My hair color is dirty blonde.<br>B: Hmm. Very interesting. Were you homeschooled? I was. You sing?<br>A: I do not sing and I was not homeschooled.  |       |
| UL-TS   | Oh. I am a singer. You dance? I do. All the time. You?  | 0.30  |
| CT      | Oh. I am a singer. You dance? I do. You dance? I do.  | 0.45  |

Table 8: Examples from the open-domain dialogue task.

137 27 - December 1, 2017 - Volume 1: Long Papers, pages 986–995. Asian Federation of Natural  
138 Language Processing.

139 [4] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019.  
140 Language models are unsupervised multitask learners.

141 [5] Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards em-  
142 pathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings*  
143 *of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence,*  
144 *Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5370–5381. Association for  
145 Computational Linguistics.

- 146 [6] Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu,  
147 Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. Recipes for building  
148 an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of*  
149 *the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 -*  
150 *23, 2021*, pages 300–325. Association for Computational Linguistics.
- 151 [7] Eric Michael Smith, Mary Williamson, Kurt Shuster, Jason Weston, and Y-Lan Boureau. 2020.  
152 Can you put it all together: Evaluating conversational agents’ ability to blend skills. In *Pro-*  
153 *ceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL*  
154 *2020, Online, July 5-10, 2020*, pages 2021–2030. Association for Computational Linguistics.
- 155 [8] Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. 2022. A  
156 contrastive framework for neural text generation. *CoRR*, abs/2202.06417.
- 157 [9] The PyTorch Lightning team. Lightning transformers. [https://github.com/](https://github.com/PyTorchLightning/lightning-transformers)  
158 [PyTorchLightning/lightning-transformers](https://github.com/PyTorchLightning/lightning-transformers).
- 159 [10] Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston.  
160 2020. Neural text generation with unlikelihood training. In *8th International Conference on*  
161 *Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenRe-  
162 view.net.
- 163 [11] Falcon William and The PyTorch Lightning team. 2019. PyTorch lightning. [https:](https://github.com/PyTorchLightning/pytorch-lightning)  
164 [//github.com/PyTorchLightning/pytorch-lightning](https://github.com/PyTorchLightning/pytorch-lightning).
- 165 [12] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony  
166 Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer,  
167 Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain  
168 Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-  
169 of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical*  
170 *Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos,*  
171 *Online, November 16-20, 2020*, pages 38–45. Association for Computational Linguistics.
- 172 [13] Omry Yadan. 2019. Hydra - A framework for elegantly configuring complex applications.  
173 <https://github.com/facebookresearch/hydra>.
- 174 [14] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston.  
175 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings*  
176 *of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Mel-*  
177 *bourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2204–2213. Association  
178 for Computational Linguistics.