

# Chat-Ghosting: Methods for Auto-Completion in Dialog Systems

Anubhab Mandal<sup>1</sup>, Sandeep Mishra<sup>1</sup>, Bishal Santra<sup>2</sup>, Tushar Abhishek<sup>2</sup>,  
Pawan Goyal<sup>1</sup>, Manish Gupta<sup>2</sup>

<sup>1</sup>IIT Kharagpur, India <sup>2</sup>Microsoft, India

anubhab.saie@gmail.com, sandeepmishraismyname@gmail.com, bishalsantra@microsoft.com  
tabhishek@microsoft.com, pawang@cse.iitkgp.ac.in, gmanish@microsoft.com

## Abstract

Ghosting, the ability to predict a user’s intended text input for inline query auto-completion, is an invaluable feature for modern search engines and chat interfaces, greatly enhancing user experience. By suggesting completions to incomplete queries (or prefixes), ghosting aids users with slow typing speeds, disabilities, or limited language proficiency. Ghosting is a challenging problem and has become more important with the ubiquitousness of chat-based systems like ChatGPT, Copilot, etc. Despite the increasing prominence of chat-based systems utilizing ghosting, this challenging problem of Chat-Ghosting has received little attention from the NLP/ML research community. There is a lack of standardized benchmarks and relative performance analysis of deep learning and non-deep learning methods. We address this through an open and thorough study of this problem using four publicly available dialog datasets: two human-human (DailyDialog and DSTC7-Ubuntu) and two human-bot (Open Assistant and ShareGPT). We experiment with various existing query auto-completion methods (using tries), n-gram methods and deep learning methods, with and without dialog context. We also propose a novel entropy-based dynamic early stopping strategy. Our analysis finds that statistical n-gram models and tries outperform deep learning based models in terms of both model performance and inference efficiency for seen prefixes. For unseen queries, neural models like T5 and Phi-2 lead to better results. Adding conversational context leads to significant improvements in ghosting quality, especially for Open-Assistant and ShareGPT. We make code and data publicly available<sup>1</sup>.

## 1 Introduction

AI-driven dialog systems or “chatbots” have become more important than ever before with the

<sup>1</sup><https://github.com/blitzprecision/Chat-Ghosting>

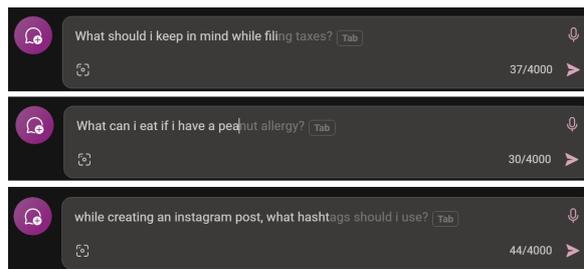


Figure 1: Ghosting Examples from Microsoft Copilot

advent of LLM-based web interfaces like ChatGPT, Copilot, Gemini, etc. Here people can interact with models like GPT-4 and Gemini-Pro on both open-ended and domain-specific topics in the form of text-based chats. With the increasing popularity of such systems, it becomes imperative to make it easier for people to interact with these agents.

One way to ease such interactions is via ghosting, the focus of this paper. Ghosting is a type-ahead completion task that predicts a user’s intended input for inline query auto-completion (QAC). Given a partially typed query (the prefix), ghosting provides a single suggestion, unlike QAC, which typically shows  $\sim 10$ . Fig. 1 shows ghosting examples from Microsoft Copilot. Like QAC, ghosting helps users save time and, at times, offers creative suggestions that enable more expressive queries.

Ghosting is a challenging task. While being effective with short ambiguous prefixes, a ghosting system also needs to be efficient. It should consider the context of the user input to generate relevant suggestions. Context could be noisy and tricky to incorporate for n-gram based models. Processing long contexts could also result into higher latency.

Although there is a vast literature on next utterance prediction using dialog models (Zhang et al., 2020; Roller et al., 2021; Adiwardana et al., 2020), the problem of completing the current utterance as the user types her prefix (i.e., ghosting), has been under-studied (Ramachandran and Murthy,

2019; Chen et al., 2019; Trajanovski et al., 2021; Goren et al., 2024). Hence, in this paper, we aim to properly define and establish a foundation for the chat-ghosting task. This includes setting up open standards for datasets and investigating effectiveness of existing language models for the task.

One of the major challenges for chat-ghosting is the lack of benchmark datasets. Hence, we adapt dialog response generation datasets for ghosting. Specifically, we curate two human-human conversation datasets, DailyDialog (DD) (Li et al., 2017) and DSTC7-Ubuntu (DU) (D’Haro et al., 2020), to model both open-ended conversations and domain-specific conversations, respectively. Further, we leverage two human-bot datasets: Open Assistant (OASST) (Köpf et al., 2024) and ShareGPT (SGPT)<sup>2</sup> with long bot utterances.

Further, in this paper, we investigate several approaches to solve the chat-ghosting task, including QAC methods, n-gram models, neural language models, and prompt engineering. The most widely used QAC method is the most popular completion (MPC) supported by tries built on historical query logs. Query Blazer (QB) (Kang et al., 2021) is an n-gram method that consists of a subword unit encoder and a classical n-gram language model to generate query completions. In addition, we also fine-tune Transformer-based language models such as GPT2 (Brown et al., 2020) and a sequence-to-sequence model such as T5 (Roberts et al., 2019). Lastly, we conduct experiments with prompting-based models such as Phi-2 (Gunasekar et al., 2023), Mistral-7B (Jiang et al., 2023) and GPT4 (Achiam et al., 2023). We experiment with these models under the context-based setting (where we use previous few utterances in the conversation as supporting context) and without context. For QAC methods and n-gram-based models, we use context for reranking. But for neural language models and prompt engineering-based methods, we condition the generation on the context by prepending it to the prefix.

Although ghosting is a natural language generation (NLG) task, standard NLG metrics like BLEU, ROUGE, CIDEr etc. are not suitable for this task as we need to rely on *prefix-match* instead of n-gram or semantic matches. We measure output quality based on exact or partial match of the prediction with respect to the ground truth, using met-

rics like Match rate (MR), Partial Recall (*P-Rec*), Partial Precision (*P-Prec*). Further, the tradeoff between the trigger rate (TR) and these quality metrics can vary significantly across methods. Hence, for fair comparison across methods, we capture quality metrics at *maximum* TR for our main results. That said, we also show the tradeoff by varying TR for various methods. We measure the number of keystrokes the user is saved from typing using typing effort saved (TES) inspired by metrics defined in Trajanovski et al. (2021). For completeness, we also present evaluation using METEOR and using GPT-4o judgments. Since low latency is crucial for ghosting, we also report the same for various methods.

We make the following main contributions: (1) To the best of our knowledge, this is the first comprehensive work on ghosting for dialog/chat bot applications, providing a novel perspective on text prediction. (2) We investigate a diverse set of methodologies for this task both with and without context: trie-based most popular completion, n-gram-based QueryBlazer (QB), Transformer-based finetuned models like T5 and GPT2, and prompt-engineering methods like Phi-2, Mistral-7B and GPT4. (3) We introduce an entropy-based dynamic early stopping strategy that adaptively halts generation when model confidence drops, yielding more precise and efficient suggestions compared to static truncation. (4) Through a comprehensive comparative study on prominent datasets such as DD (open-domain), DU (task-specific), OASST and SGPT via multiple evaluation metrics, we compare the generation quality and latency of chat-ghosting across different methods, shedding light on their performance dynamics and computational efficiency.

## 2 Related Work

### 2.1 Query Auto-Completion

Bast and Weber (2006) first introduced the efficacy of QAC in search engines. Ranking for QAC systems is typically supported by several modules such as most popular completion (Whiting and Jose, 2013), time sensitive suggestions (Shokouhi and Radinsky, 2012; Wang et al., 2017), location sensitive suggestions (Backstrom et al., 2008; Welch and Cho, 2008), ghosting (Ramachandran and Murthy, 2019), session co-occurrences (Bar-Yossef and Kraus, 2011) and non-prefix matches (Gog et al., 2020). Besides these, suggestion ranking in QAC has also been studied using traditional Ma-

<sup>2</sup><https://sharegpt.com/>, dataset version that was used: anon8231489123ShareGPT\_Vicuna\_unfiltered

	DD Train	DD Test	DU Train	DU Test	OASST Train	OASST Test	SGPT Train	SGPT Test
Avg. words per utterance	12.37	12.42	13.48	14.55	20.36	23.35	53.27	56.85
Avg. characters per utterance	57.05	57.33	73.96	79.57	115.05	131.85	341.01	366.88
No. of utterances	69,216	7,986	549,002	5,588	19421	981	328078	1088

Table 1: Statistics of the all 4 Datasets - DD, DU, OASST & SGPT

chine Learning (ML) methods (Di Santo et al., 2015; Jiang et al., 2014; Sordoni et al., 2015) and deep learning (DL) methods like convolutional latent semantic model (Mitra and Craswell, 2015), LSTMs (Wang et al., 2020) and BERT and BART (Mustar et al., 2020). Personalization has also been studied for QAC systems (Shokouhi, 2013; Song et al., 2017; Fiorini and Lu, 2018; Yin et al., 2020). Recently, multimodal auto-completions have also been explored (Mishra et al., 2026). There have also been attempts to *generate* effective QAC suggestions (Park and Chiba, 2017; Kim, 2019; Jiang and Wang, 2018; Sordoni et al., 2015; Lee et al., 2021; Olteanu et al., 2020; Maurya et al., 2023). In this work, we focus on the chat-ghosting task. Unlike showing multiple full completions in the QAC systems, ghosting aims at generating a single relevant suggestion with next few tokens.

## 2.2 Dialog Response Generation

Several Sequence-to-Sequence models (Ritter et al., 2011; Santra et al., 2021a), latent variable models (Zhao et al., 2017; Bao et al., 2019), pretrained models (Zhang et al., 2020; Roller et al., 2021; Adiwardana et al., 2020), response retrieval/next-utterance selection models (Lowe et al., 2015; Chen and Wang, 2019; Whang et al., 2021; Xu et al., 2021; Humeau et al., 2019; Henderson et al., 2020) and retrieval augmented generation (Cai et al., 2021; Komeili et al., 2021) models have been proposed for dialog response generation. In this work, we propose the task of ghosting for dialog systems. While dialog generation aims to generate full utterances, ghosting aims to generate next few tokens given a few characters already typed by the user. We discuss detailed related work in Appendix A.

## 3 Chat-Ghosting: Task and Datasets

Ghosting aims to suggest accurate query completions based on an incomplete, partial prefix.

### 3.1 Task Formulation

The input to a chat-ghosting model comprises a partially written user prefix ( $p$ ) and optional dialog history or context ( $h$ ). The model is expected

to predict a suitable completion ( $c$ ) for the prefix  $p$  such that  $[p; c]$  is a valid response for dialog context  $h$ . Thus, at train time, the aim is to learn parameters  $\theta$  of the chat-ghosting model as  $\theta^* = \operatorname{argmax}_{\theta} P(c|p, h, \theta)$ . The trained model can then be used to generate predictions using  $\hat{c} = \operatorname{argmax}_c P(c|p, h, \theta^*)$ .

### 3.2 Dataset Design

There are no public datasets for training chat-ghosting models, complicating both training and evaluation. To address this, we simulate user interactions with auto-completion systems using existing dialog response generation datasets and emulate chat-ghosting outcomes.

Users typically engage in open-ended or domain-specific conversations with chatbots, making it necessary to train and evaluate the performance of the ghosting methods on both kinds of data. Hence, we use DailyDialog (DD), an open domain dialog dataset, and DSTC7-Ubuntu (DU), a domain-specific conversation dataset. DD contains open-domain English conversations, and we use a cleaned version to avoid data leaks (Wen et al., 2022). DU includes domain-specific English conversations from a Linux IRC channel.

But DD and DU include human-human conversations. Hence, we also experiment with two other datasets: OASST and SGPT which include human-bot conversations. OASST is a human-annotated assistant conversation corpus. SGPT contains user-LLM-chatbots conversations collected by the ShareGPT API. Table 1 presents the basic statistics of these datasets. Since the main aim of ghosting is to assist a human, for both OASST and SGPT, we just consider just the human utterances for both training and inference to measure the effectiveness.

## 4 Methods for Chat-Ghosting

We follow various modeling strategies for chat-ghosting that gives a complete view at the spectrum of trade-off between accuracy, latency and coverage. We investigate several ML and DL approaches for the task, including QAC methods, n-gram models, neural language models and prompt engineer-

ing. We experiment with both the “with context” and “without context” settings.

#### 4.1 Standard QAC Methods

We experiment with two standard QAC methods: MPC and MPC++. Given a prefix, the MPC model, proposed in [Bar-Yossef and Kraus \(2011\)](#), extracts a limited number ( $k$ ) of completions from a character-trie structure (also referred to as the main trie) built using a corpus of past user-utterances. Prefixes for which completions are available through MPC are termed as ‘seen prefixes’, while those lacking completions are termed as ‘unseen prefixes’.

Main trie provides good coverage for QAC, however, for dialog systems, it is less likely that incoming user utterances would match with those in the training set. Hence, the coverage of the system can be very low due to novel utterances. To improve the coverage, we design an alternative method that utilizes a suffix-trie to handle unseen prefixes. The suffix trie is built using word n-gram suffixes (with  $\text{freq} \geq 2$ ) of the training utterances. Both the main trie and the suffix tries are character-level tries. The two tries can also be used in combination. Thus, in the MPC++ method, seen prefixes utilize the main trie, whereas unseen prefixes utilize the suffix trie.

#### 4.2 N-Gram Model: Query Blazer (QB)

QB is a fully-generative QAC system that uses a longest prefix match (LPM) subword unit encoder and a subword-level n-gram language model to generate query completions. QB can suggest queries from log history and generate new words or phrases. By leveraging the n-gram model’s data structure, QB materializes partial completions offline, reducing runtime overhead.

During training, QB extracts a subword vocabulary from the dataset and constructs an encoder as a finite state transducer (FST). The trained language model is represented as a weighted FST. During inference, based on the input prefix, the language model performs beam search to generate the final top- $k$  completions. The confidence score for a prediction is the sum of the negative log-likelihoods of the predicted tokens, normalized by prediction length. This confidence score is used as a threshold for triggering decisions.

#### 4.3 Neural Language Modeling Methods

We also experiment with popular Transformer-based models for generating ghosting suggestions.

Specifically, we use the GPT2-base ([Brown et al., 2020](#)) (decoder-only) and T5-base ([Roberts et al., 2019](#)) (Seq2Seq) models. GPT2-base has 12 decoder layers, while the T5-base has 12 encoder and 12 decoder layers. Our objective is to express the ghosting task as a causal language modeling task and investigate the efficacy of these models.

The prefix and the completion constitute the source and target sequences, respectively, for these models. The input prefix may end with a complete word, with a space after a complete word, or in the middle of an incomplete word. For example, for the query “How are you”, consider the three cases where “prefix|suffix” pairs are (1) “How are|you”, (2) “How are |you”, (3) “How ar|e you”. In the first case, the model needs to learn to predict a space first while in the second and third cases, the model must generate a valid character. However, T5 and GPT2 tokenizers differ in the way they handle spaces, requiring us to handle the three cases carefully as discussed in [Appendix B](#).

#### 4.4 Prompt Engineering Methods

We experiment with Large language models (LLMs) like GPT-4 ([Achiam et al., 2023](#)), and small language models (SLMs) like Phi-2 ([Gunasekar et al., 2023](#)) and Mistral-7B ([Jiang et al., 2023](#)) in zero-shot mode. We also finetune (FT) Phi-2 on our training datasets. Our experiments show that we could obtain best results by prompting the models to generate a max of 2 words for Mistral, 10 words for Phi-2 (PT) and 3 words for Phi-2 (FT). In the zero-shot mode, we prompt ([Appendix D](#)) the models to generate appropriate completions given a prefix.

#### 4.5 Contextual Ghosting

In chatbots, one can expect ghosting quality to significantly improve by considering previous utterances in the current conversation as additional context. Thus, we modify the methods described so far in this section to make use of previous user utterances while generating ghosting suggestions.

For neural language modeling methods and prompt engineering models, we prepend the context to the current user prefix. For prompt engineering-based methods, the detailed prompt is in [Appendix E](#). To incorporate historical context into QAC and n-gram models (QB, MPC, MPC++), we use a reranking setup. For each prefix in the test set, we first retrieve the top- $k$  completions ( $k = 10$ ) and their corresponding

scores. Next, we calculate the cosine similarity between each (prefix + completion) pair and the historical context. This cosine similarity is computed between the TF-IDF vectors of the text pairs, using a TF-IDF vectorizer trained on the training set to capture a dataset-specific vocabulary. Score for each completion is then calculated as  $\alpha \cdot (\text{original\_score\_scaled}) + \beta \cdot (\text{cos\_sim}) + \gamma \cdot (\text{length\_penalty\_scaled})$ , where each term is scaled to  $(-1, 1)$  for consistency across components. The hyperparameters  $\alpha$ ,  $\beta$ , and  $\gamma$  are tuned on a validation set, with  $(\alpha, \beta, \gamma) = (0.5, 0.3, 0.2)$  yielding the best results. Length penalty is calculated as  $\frac{1}{1 + \text{completion\_length}}$ . It helps to reduce the bias towards longer completions, which tend to have higher cosine similarity due to more matching words. The length penalty helps prioritize shorter, more relevant completions.

#### 4.6 Ghosting Evaluation Metrics

Standard NLG metrics like BLEU, ROUGE, etc. are not suitable for ghosting as we need to rely on prefix-match instead of n-gram or semantic matches. Any deviation from prefix matches typically leads to higher cognitive load in examining the suggestion, leading to less likelihood of suggestion acceptance. Hence, we design new metrics suitable for ghosting.

Quality Metrics: When a ghosting suggestion is shown, the suggestion text may match exactly with the ground truth or match just an initial part of the target. Hence, we define Match Rate (MR) as the ratio of number of times the suggestion exactly matched the ground truth to the number of times suggestions were shown, similar to [Trajanovski et al. \(2021\)](#). For real world usability, MR is too strict as users may accept partially correct suggestions and make the necessary minor adjustments. Further, users may show some bias towards correctness on the prefix side of the prediction than near the end. Hence, to allow for partial matches, we define Partial Recall (*P-Rec*) and Partial Precision (*P-Prec*). *P-Prec* (or *P-Rec*) is defined as the fraction of the total character length of the prediction (or ground truth) by which the prefix of prediction and ground truth overlap. For completeness, we also present evaluation using METEOR and using GPT-4o judgments in [Appendix M](#).

Trigger Rate (TR): At every character, the ghosting system can show a suggestion or not based on the confidence score of the suggestion. For users with high typing speed, intent clarity and good language

skills, frequent ghosting may be distracting. On the other hand, for users who are slow at typing or who struggle in query formulation, triggering ghosting at every character may be acceptable. Although defining a threshold based on user traits is out of scope of the current paper, we introduce TR as an operating point. It is the ratio of number of times a ghosting suggestion was shown to the number of characters typed by the user.

If  $TR < 100\%$  (or whatever is max possible for a model), the probability of low confidence predictions will decrease and the system will suggest completions only for some subset of all the test cases. Hence, for fair comparison of various modeling approaches, we compare their performances with respect to complete match (MR) or partial match (*P-Rec* and *P-Prec*) metrics at max TR.

Typing Effort Saved (TES): Given an utterance in the test set, we first list down all possible (prefix, completion) splits. Thus, an utterance of length  $l$  leads to  $l - 1$  (prefix, completion) samples. Metrics like MR, *P-Rec* and *P-Prec* are computed per prefix and micro-averaged across the entire test set. On the other hand, the TES metric is computed per utterance in a manner such that it simulates the interaction of the user with the auto-completion system, based on a greedy selection heuristic, until the entire utterance is typed out. At any character position in the utterance, the auto-completion system generates a completion and we assume that the user either (a) accepts it if the full suggestion of length  $s$  matches exactly with the first  $s$  characters of the target completion, or (b) rejects the suggestion if it does not match, and types the next character. This process continues until the full ground truth utterance has been generated. At the end of the process, the TES metric is computed as 1 minus the ratio of the number of characters typed by the user to the utterance length. For example, let the utterance be “who am I?”. Let the process be as follows: user types “w”, user accepts “ho”, user types “[space]”, user does not accept “is” and continues to type “a”, user accepts “m i?”. Here the user typed 3 characters and utterance length is 9. Hence,  $TES = 1 - (3/9) = 2/3$ . Note that TES is closely related to Est. ChS ([Trajanovski et al., 2021](#)) and eSaved ([Kharitonov et al., 2013](#)).

## 5 Experiments and Results

We aim to determine which model performs the best for the chat-ghosting task in with/without and

Models	Unseen (94%)						Seen (6%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	MLen	MR	P-Rec	P-Prec	TES	Pred Len	MLen	MR	P-Rec	P-Prec	TES	Pred Len	MLen
<b>QB</b>	5.87	10.72	<b>43.31</b>	<b>31.63</b>	4.1	1.4	26.10	34.81	57.69	41.92	4.5	2.3	6.28	11.21	<b>43.60</b>	<b>32.24</b>	4.1	1.4
<b>MPC</b>	0.00	7.54	19.47	2.90	<b>29.8</b>	2.1	<b>48.92</b>	<b>61.10</b>	<b>62.77</b>	<b>46.13</b>	14.4	<b>7.5</b>	4.30	11.98	23.06	5.50	<b>28.5</b>	<b>2.6</b>
<b>MPC++</b>	5.14	11.77	28.49	13.85	19.2	1.8	<b>48.92</b>	<b>61.10</b>	<b>62.77</b>	<b>46.13</b>	14.4	<b>7.5</b>	6.41	13.20	29.49	15.79	19.1	2.0
<b>T5 (220M)</b>	<b>6.10</b>	13.86	25.16	20.88	18.9	2.2	32.07	44.30	49.70	44.07	10.7	3.5	<b>6.61</b>	14.45	25.64	22.27	18.7	2.3
<b>GPT2 (124M)</b>	5.26	11.66	27.64	24.47	12.3	1.9	28.50	39.14	50.61	45.30	8.0	2.9	5.72	12.20	28.10	25.72	12.3	1.9
<b>Phi-2 (2.7B) (PT)</b>	2.82	6.00	17.21	12.38	10.7	0.9	15.01	25.24	34.72	22.84	10.6	1.8	3.10	6.44	17.61	13.01	10.7	0.9
<b>Phi-2 (2.7B) (FT)</b>	0.90	<b>15.59</b>	27.54	13.12	9.0	<b>2.5</b>	3.52	38.59	33.12	15.07	8.5	2.8	0.95	<b>16.04</b>	27.65	13.23	9.0	2.5
<b>Mistral-7B (PT)</b>	1.07	6.03	8.04	5.74	17.1	0.9	6.04	24.17	16.27	10.43	<b>20.1</b>	1.7	1.17	6.40	8.21	6.02	17.1	0.9
<b>GPT4</b>	3.16	13.81	18.52	12.11	18.0	2.3	4.98	33.00	20.77	10.82	18.8	2.6	3.20	14.19	18.56	12.03	18.0	2.3

Table 2: Results of various approaches on the Chat-Ghosting task for the DD dataset. PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as shown in Table 4. Pred Len and Matched Len (MLen) are in chars.

Models	Unseen (48.2%)						Seen (51.8%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	MLen	MR	P-Rec	P-Prec	TES	Pred Len	MLen	MR	P-Rec	P-Prec	TES	Pred Len	MLen
<b>QB</b>	2.31	8.05	<b>39.77</b>	<b>24.70</b>	5.3	1.7	6.22	12.26	47.46	32.43	6.1	2.7	4.29	10.19	43.66	28.70	5.7	2.2
<b>MPC</b>	0.00	5.73	15.34	2.01	<b>53.8</b>	2.0	<b>86.49</b>	<b>87.37</b>	<b>88.81</b>	<b>74.07</b>	<b>71.0</b>	<b>64.5</b>	<b>75.07</b>	<b>76.57</b>	<b>79.09</b>	39.24	<b>68.7</b>	<b>56.2</b>
<b>MPC++</b>	2.03	8.44	26.82	10.42	32.4	1.9	<b>86.49</b>	<b>87.37</b>	<b>88.81</b>	<b>74.07</b>	<b>71.0</b>	<b>64.5</b>	54.22	57.21	65.13	<b>43.30</b>	56.2	40.6
<b>T5 (220M)</b>	<b>3.93</b>	<b>11.76</b>	24.59	16.82	26.1	<b>2.9</b>	5.86	13.51	26.07	21.61	27.5	3.4	4.90	12.64	25.34	19.30	26.8	3.2
<b>GPT2 (124M)</b>	3.18	9.98	25.08	19.02	20.3	2.4	7.84	14.54	30.06	27.26	22.3	4.1	5.53	12.28	27.60	23.28	21.3	3.3
<b>Phi-2 (2.7B) (PT)</b>	0.76	4.59	19.37	9.81	13.1	1.1	1.00	5.06	20.42	10.73	13.3	1.2	0.88	4.83	19.90	10.29	13.2	1.2
<b>Phi-2 (2.7B) (FT)</b>	0.67	11.08	31.92	13.37	8.3	2.5	0.80	11.38	32.11	13.74	8.2	2.5	0.74	11.23	32.02	13.56	8.3	2.5
<b>Mistral-7B (PT)</b>	0.25	3.78	7.51	3.90	18.5	1.0	0.28	3.99	7.63	3.75	18.4	1.0	0.26	3.88	7.57	3.82	18.5	1.0
<b>GPT4</b>	1.40	9.21	16.55	9.55	19.7	2.2	1.37	9.29	16.55	9.08	19.9	2.3	1.38	9.25	16.55	9.31	19.8	2.3

Table 3: Results of various approaches on the Chat-Ghosting task for the DU dataset. PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as shown in Table 4. Pred Len and Matched Len (MLen) are in chars.

	DD			DU		
	Unseen	Seen	Full	Unseen	Seen	Full
<b>QB</b>	86.28	90.61	86.37	85.81	86.60	86.21
<b>MPC</b>	21.92	100.00	23.44	15.51	100.00	58.11
<b>MPC++</b>	66.25	100.00	66.91	62.88	100.00	81.60
<b>T5 (220M)</b>	100.00	100.00	100.00	99.98	99.99	99.98
<b>GPT2 (124M)</b>	95.41	97.57	95.45	96.46	96.87	96.67
<b>Phi-2 (2.7B) (PT)</b>	59.14	70.08	59.35	49.73	48.92	49.32
<b>Phi-2 (2.7B) (FT)</b>	99.49	99.83	99.50	90.33	91.69	91.02
<b>Mistral-7B (PT)</b>	65.58	68.64	65.64	55.54	54.66	55.10
<b>GPT4</b>	98.31	99.57	98.34	90.24	91.58	90.92

Table 4: Max TR values. Values<100% imply that models generate empty predictions.

human-human vs human-bot settings. Additionally, we explore the trade-offs between model performance and coverage. Detailed hyper-parameter settings can be found in Appendix F.

### 5.1 Main Ghosting Results (without context)

Tables 2 and 3 show accuracy metrics and prediction lengths for unseen/seen/full test sets across various models on DD and DU, resp. Due to lack of space, we report results for OASST and SGPT in Tables 9 and 10 resp. in Appendix C. Here “seen/unseen” implies that the utterance was seen/unseen in the train set. These results have been computed at max TR, i.e., no thresholding of confidence values was done to compute these results, and thus comparison is made with respect to all samples in the test sets. However, as shown in Table 4, the max TR may not be 100% for all models as models may generate empty predictions for some queries.

From Tables 2 and 3, several key observations

emerge. For unseen queries, the QB model proves to be the most effective for ghosting, achieving the highest TES and *P-Prec* scores, followed by the MPC++, finetuned T5, and GPT-2 models. Specifically, QB has the highest *P-Prec* (~43 in DD and ~40 in DU), indicating that on average, the first 42% characters of its predictions are correct. However, QB’s *P-Rec* remains low. Despite QB’s higher *P-Prec* compared to T5 or GPT-2, its individual completions are much shorter, averaging 4.1 and 5.3 characters for DD and DU, respectively. While T5 and GPT-2 may have longer prefix matches in their suggested completions, longer incorrect completions could negatively impact user experience.

When comparing “trained” language models (LMs) to other solutions like trie and prompt engineering, “trained” LMs emerge as the clear winners. Within this family, the classical frequency-based LM QB, likely due to its implicit memorization capabilities, outperforms the trained neural LMs such as GPT-2, T5, and Phi-2 (FT). As expected, trie-based solutions, which rely solely on memory, significantly outperform other models in the seen test set. MPC++, which combines trie with some language modeling capabilities (through the use of suffix-trie), performs better than MPC for unseen queries but falls short for seen queries. MPC also provides the longest suggestions. Considering both MR and TR, T5 is the best model for unseen queries, while finetuned Phi-2 outperforms pretrained Phi-2 except on MR.

	DD												DU											
	Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars			Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars		
	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec
MPC	0.01	4.61	22.93	0.06	7.33	17.82	0.65	11.71	15.82	8.17	<b>29.56</b>	29.78	0.04	3.10	20.00	0.19	6.14	12.71	0.33	8.46	11.06	0.79	<b>14.75</b>	17.81
MPC++	0.01	4.61	22.93	0.39	7.48	18.86	4.05	12.08	26.96	<b>9.06</b>	15.91	34.68	0.04	3.10	20.00	0.25	6.08	14.15	1.59	8.06	23.28	2.44	9.07	30.92
QB	<b>0.13</b>	3.48	<b>36.09</b>	1.14	5.49	<b>42.38</b>	4.80	9.88	<b>44.38</b>	8.30	13.52	<b>44.80</b>	0.13	3.29	<b>29.05</b>	0.95	6.24	<b>37.69</b>	2.07	8.01	<b>43.38</b>	2.77	8.83	<b>42.46</b>
T5	0.08	<b>5.11</b>	22.36	<b>1.21</b>	<b>8.34</b>	21.97	<b>5.31</b>	<b>13.40</b>	24.76	8.55	16.68	27.65	<b>0.14</b>	<b>3.83</b>	20.67	<b>0.96</b>	<b>7.00</b>	22.22	<b>2.83</b>	<b>9.99</b>	24.90	<b>4.16</b>	11.49	25.32

Table 5: Results of various approaches on the Chat-Ghosting task, on the DD and DU dataset for different prefix lengths. Metrics at max TR and reported for unseen test set.

	DD		DU	
	Seen	Unseen	Seen	Unseen
<b>Prefix</b>	listen to the	which kind of salad are you goi	unable to fetch some archives, maybe run apt-get update or	so do you know h
<b>Ground Truth (GT)</b>	sound of nature ! it's like music	ng to make	try with -fix-missing?	owto setup a static route in the routing table
<b>QB</b>	radio	ng	try with -fix-missing?	ow?
<b>MPC</b>	sound of nature ! it's like music	[EMPTY]	try with -fix-missing?	ow?
<b>MPC++</b>	sound of nature ! it's like music	ng	try with -fix-missing?	ow?
<b>T5 (220M)</b>	music, please	ng to make	something?	ow to do that?
<b>GPT2 (124M)</b>	news	ng to make	acle-java7-jre will help	ow to do that?
<b>Phi-2 (2.7B) (PT)</b>	music	ng to make	check for updates	ow to use a computer
<b>Phi-2 (2.7B) (FT)</b>	music. it	ng to have	something?	ow to do
<b>Mistral-7B (PT)</b>	radio station for classic hits	ng to make?	try using a different package manager	ow to make a paper airplane?
<b>GPT4</b>	latest podcast	ng to make?	try a different mirror	ow to fix this problem?

Table 6: Examples of suffixes predicted by various non-contextual models for different prefixes. More examples from both contextual and non-contextual models are in Appendix G.

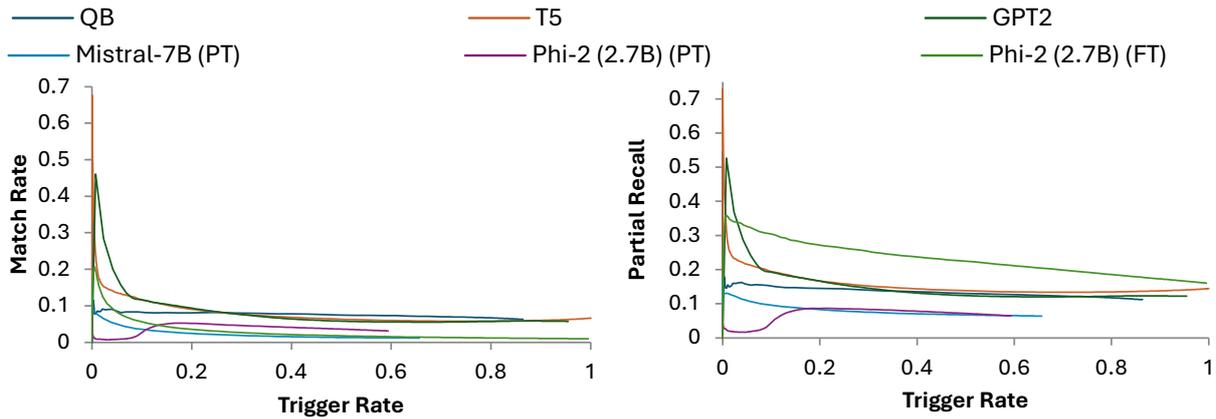


Figure 2: MR for DD

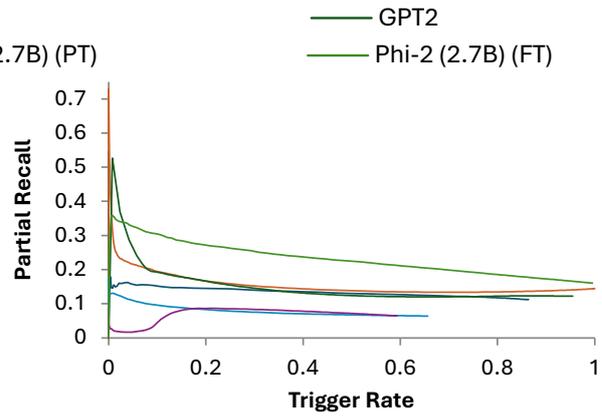


Figure 3: P-Rec for DD

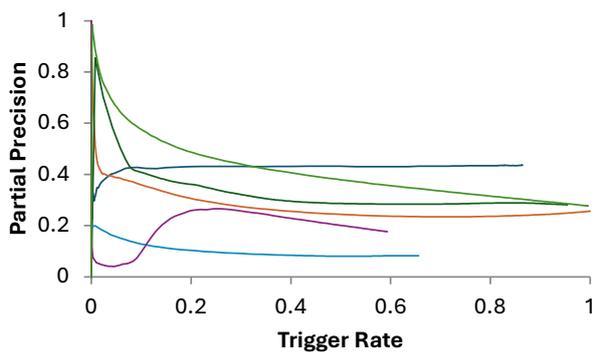


Figure 4: P-Prec for DD

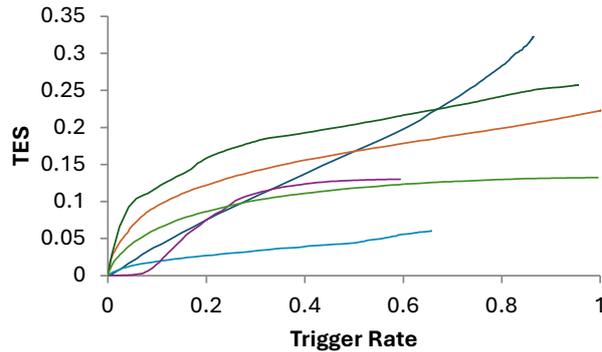


Figure 5: TES for DD

Overall, these results suggest that for effective ghosting, it is beneficial to use LMs with inherent memory capacity and train them on logs. If a query seen during training is encountered during testing, the model reiterating the exact same completion

is preferred over a model hallucination. Models that cannot be trained, such as zero-shot LLMs, are not effective for ghosting. It is important to note that all models perform poorly on the unseen test set, highlighting the challenging nature of chat-

ghosting. Transformer-based LMs, regardless of their scale, also fail to perform satisfactorily. Zero-shot, instruction-tuned SLMs and LLMs perform significantly worse than their finetuned counterparts, underscoring the necessity of training with available logs for reliable performance. Based on our findings, we recommend using MPC for seen prefixes, and T5 or QB can be employed for unseen prefixes, depending on the latency requirements.

For OASST and SGPT, where utterance lengths are much longer, we observe that the pretrained knowledge helps and Mistral and Phi-2 provide the best *P-Rec* and MR resp for the unseen set. Phi-2 also provides the best TES on unseen set. On the seen sets, MPC, MPC++ and QB continue to perform better even on these human-bot datasets.

## 5.2 Truncation-based results

Various methods generate predictions of different sizes. Intuitively, longer suggestions have a lesser probability of being correct. This implies that MR and TES should be higher for shorter suggestions. To verify this we truncate the predicted suggestions to  $t$  words where we vary  $t$  from 1 to 10. We show results for MR, *P-Prec*, *P-Rec* and TES with DD and DU datasets in contextual and non-contextual settings across all the methods in Figs. 10 to 17 in Appendix K for the unseen test set. GPT2 and T5 provide high TES (better than QB), when truncating to 1-2 words.

## 5.3 Dynamic early stopping

Static truncation of outputs to a fixed number of tokens may not be optimal across all test samples. A more adaptive approach is to halt generation when model confidence in the next token drops. This confidence can be estimated using entropy of the token probability distribution.

We apply entropy-based dynamic early stopping with thresholds of 0.6 and 3 during suggestion generation using T5 and GPT2 on the unseen test sets of DD and DU in the non-contextual setting. As shown in Table 7, a threshold of 0.6 significantly boosts *P-Prec*, with T5 reaching 43.95 and GPT2 46.16 on the DD dataset. This improvement comes with shorter predictions, indicating more concise outputs, and a marked increase in TES, reflecting reduced keystrokes. Overall, dynamic early stopping offers a better trade-off across metrics compared to static truncation.

Model	Dataset	Threshold	TR	MR	P-Prec	P-Rec	Pred Len	TES
T5	DD	None	100.00	6.10	25.16	13.86	18.9	20.88
T5	DD	3	96.42	5.80	33.14	12.99	14.5	29.77
T5	DD	0.6	88.03	5.56	43.95	12.50	7.0	45.80
T5	DU	None	99.98	3.93	24.59	11.76	26.1	16.82
T5	DU	3	95.92	3.71	33.81	11.06	21.0	27.99
T5	DU	0.6	91.62	2.94	48.97	9.19	11.6	44.41
GPT2	DD	None	95.41	5.26	27.64	11.66	12.3	24.47
GPT2	DD	3	97.56	4.61	35.83	10.41	9.2	34.75
GPT2	DD	0.6	99.38	3.81	46.16	8.71	4.1	45.83
GPT2	DU	None	96.46	3.18	25.08	9.98	20.3	19.02
GPT2	DU	3	97.71	2.73	34.29	8.95	17.8	31.44
GPT2	DU	0.6	98.82	2.02	47.01	7.25	7.3	45.72

Table 7: Entropy-based Dynamic Early Stopping

## 5.4 Contextual Ghosting Results

Tables 14 and 15 in Appendix G show accuracy metrics and prediction lengths for unseen/seen/full test sets across various models on OASST and SGPT datasets, respectively for the “with context” setting. Table 16 in Appendix G shows max TR values for the “with context” setting. From Tables 14 and 15, we observe that (1) Context significantly improves the accuracy of T5 for the unseen test samples. (2) Adding context does not lead to any significant improvements across most models except T5. This could be because of noise in context or irrelevance to current utterance.

Similarly, Tables 11 and 12 in Appendix C show accuracy metrics and prediction lengths for unseen/seen/full test sets across various models on DD and DU datasets, respectively for the “with context” setting. Table 13 in Appendix C shows max TR values for the “with context” setting. We observe that context significantly improves the best accuracy across all metrics for the unseen test samples. Also, context helps improve neural models much more compared to QB, MPC or MPC++ where context is used only for reranking.

## 5.5 Prefix Bucketing Results

To benchmark the effectiveness of various models on different prefix lengths, we analyze their performance on 4 prefix length buckets (1-5, 6-12, 13-25 and 26-50 characters). Table 5 shows the results from specific prefix length buckets on the unseen splits for the best 4 models<sup>3</sup> of the test set. For DD, the overall test set contains 36k, 50k, 86.3k and 116k samples across the four buckets. For DU, there are 13k, 18k, 31.4k and 48.2k samples across the four buckets.

Different models show their strengths in different prefix buckets. Generally, as prefix length in-

<sup>3</sup>Prefix bucketing results for the seen test set are in Appendix H. Prefix bucketing results for “with context” setting are in Appendix I.

creases, MR and *P-Prec* improve. For both DD and DU datasets, MPC achieves higher *P-Prec* at longer prefixes by leveraging historical utterances, where the candidate space is more constrained. The same holds for QB, which struggles more with shorter queries due to a larger candidate set. T5 performs best on *P-Prec* in the first three buckets, while MPC leads in the long-prefix bucket. For *P-Prec*, QB outperforms other models across all buckets.

## 5.6 Ghosting accuracy at varying TR thresholds for DD and DU

In Tables 2 and 3, we showed results computed at maximum possible TR. However, in practical systems, models are never deployed without confidence level thresholds. Hence, we show variation in accuracy metrics with varying TRs in Figs. 2-5 for DD dataset. Similar plots for DU dataset are in Appendix J. For various accuracy metrics (except TES), the expected nature of the metric-vs-TR curves is that they should be decreasing. This is because higher confidence predictions are expected to be more accurate and included when TR is low (i.e. high-confidence threshold). But there are several exceptions to this: Phi-2 (PT) in the DD dataset, and Phi-2 (FT) in DU dataset show increasing trends in metric-vs-TR plots. We believe this may be because these models are pretrained for very different kinds of tasks using cross-entropy loss which does not directly correlate with the metrics chosen here. As TES models a user that greedily chooses whenever there is a match between the model suggestion and the target, it is natural that it will always increase as more predictions are included. This is why, unlike the other metrics TES-vs-TR plot (Fig. 5 and 9) is increasing in nature. The *P-Prec* plots also show that while the finetuned Phi-2 model is not the best at max TR, it performs the best for low TR regions. Similarly, GPT2 provides best TES values for  $TR < 0.75$  although QB is best at max TR.

## 5.7 Qualitative Analysis

Table 6 shows predictions using various methods for 4 samples – one each from seen/unseen test sets of DD and DU datasets. We observe that while several of these predictions may not exactly match with the ground truth completions, they are semantically meaningful.

## 5.8 Inference Latency and other Optimizations

Table 22 in Appendix L shows inference latency for various methods. We observe that inference can be done using models like T5 and GPT2 within 100-150ms on a V100 GPU with batch size=1. Further optimizations like quantization aware training, post-training quantization and CUDA-level optimizations (as detailed in Appendix L) can be done to reduce latency and increase throughput.

## 6 Conclusion

We propose the novel problem of chat-ghosting and repurpose four public dialog datasets, DailyDialog, DSTC7-Ubuntu, Open Assistant, and ShareGPT, to benchmark trie-based QAC, ngram-based QB, deep learning, and prompt-engineered models. We adapt metrics from related tasks and evaluate model performance across multiple operating points. While trie-based QAC excels on seen queries, neural models like T5 and Phi-2, and ngram-based QB perform better on unseen ones. Our experiments with billion-parameter models show strong performance at low TR and on unseen queries in human-bot datasets, but with high inference latency. Our novel entropy-based dynamic early stopping outperforms static truncation. Overall, our findings highlight a trade-off between accuracy and latency, motivating future work on combining methods effectively.

## 7 Limitations

The work has been done for English prefixes and dialogs only. It will be nice to extend this work to queries in multiple other languages.

Experiments were performed using small encoder models for natural language generation, keeping in mind the latency constraints for query auto-completion task. We could train larger models and distill them to smaller sizes.

We analyzed results with individual methods separately. However, production systems exploit complementary nature of such methods by integrating multiple methods to decide when each of the methods will trigger. It would be nice to learn a ranking over suggestions from these multiple methods.

We experimented with a basic approach to incorporate previous utterances from dialog history. However, such context may be unrelated to current user intent and we may need to choose relevant context to improve ghosting quality further. We leave complex modeling of context like (Zhang

et al., 2020; Bao et al., 2019; Roller et al., 2021; Adiwardana et al., 2020) as future work.

This work focused on proposing the task of Chat-ghosting, and benchmarking against strong baseline solutions. We acknowledge that incorporating user traits like age, demography, typing speed, etc. is important. However, there does not exist any dialog datasets with these user traits. We hope to gather such a dataset, and build methods for personalized chat-ghosting in the future. We plan to embed the topics in historical conversations from the same user, and use them to define short term and long term user interests. Such user interest vectors can then be used for personalization by conditioning the AutoSuggest generation on such vectors.

The primary focus of this paper was to explore and evaluate methods for autocompletion of text dialogues, specifically focusing on the model’s ability to generate contextually relevant and coherent responses. Our work does not address the robustness of our model against adversarial inputs and noisy prefixes. One possible approach to handle noisy prefixes could be to predict full query completion rather than predicting just a suffix.

As a first foray into applying ghosting within dialog systems, our current work is constrained by limitations in existing datasets and the scope of initial experimentation. Lack of diverse datasets, including those that support multi-lingual and multi-modal approaches, is an important challenge. We acknowledge the potential of hybrid solutions to balance accuracy, latency, compute costs, and responsible AI considerations. Future research will focus on generating more robust datasets, exploring different hybrid combinations, and systematically evaluating their performance along these dimensions.

## 8 Ethical Considerations

Like many other pretrained language representation models, the proposed models may also have learned patterns associated with exposure bias. Interpretability associated with the output is rather limited, hence users should use the outputs carefully. The proposed model generates possible completions, and does not filter out any “problematic” completions. Thus, for applications, where candidate responses could be problematic, (e.g., offensive, hateful, abusive, etc.), users should carefully filter them out before using the output from our

model.

We leverage publicly available datasets only. These datasets do not contain any personally identifiable data including credit card numbers, email addresses, etc. as mentioned by their respective papers. Just like any other generative model, real world training and deployment of QAC models need to take care of privacy concerns and potential biases. Such models are based on past user data and therefore reflects the biases prevalent online. That said, such deployments should use several existing methods in the literature (Maheswaran et al., 2024a,b; Liu et al., 2021; Lu et al., 2022) to avoid showing offensive and biased suggestions to users.

All the datasets used in this work are publicly available. We did not collect any new dataset as part of this work. DailyDialog was downloaded from <http://yanran.li/dailydialog> and is licensed under CC BY-NC-SA 4.0. DSTC7-Ubuntu was downloaded from [https://ibm.github.io/dstc-noesis/public/data\\_description.html#ubuntu](https://ibm.github.io/dstc-noesis/public/data_description.html#ubuntu) and is available under MIT license. Open Assistant was downloaded from <https://github.com/LAION-AI/Open-Assistant> and is available under apache-2.0 license. ShareGPT was downloaded from [https://huggingface.co/datasets/anon8231489123/ShareGPT\\_Vicuna\\_unfiltered](https://huggingface.co/datasets/anon8231489123/ShareGPT_Vicuna_unfiltered) and is available under apache-2.0 license.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. 2020. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*.
- Lars Backstrom, Jon Kleinberg, Ravi Kumar, and Jasmine Novak. 2008. Spatial variation in search engine queries. In *Proceedings of the 17th international conference on World Wide Web*, pages 357–366.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2019. Plato: Pre-trained dialogue generation model with discrete latent variable. *arXiv preprint arXiv:1910.07931*.
- Ziv Bar-Yossef and Naama Kraus. 2011. Context-sensitive query auto-completion. In *Proceedings of*

- the 20th international conference on World wide web*, pages 107–116.
- Holger Bast and Ingmar Weber. 2006. Type less, find more: fast autocompletion search with a succinct index. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 364–371.
- 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.
- Hengyi Cai, Hongshen Chen, Yonghao Song, Xiaofang Zhao, and Dawei Yin. 2021. Exemplar guided neural dialogue generation. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3601–3607.
- Ciprian Chelba, Mia Chen, Ankur Bapna, and Noam Shazeer. 2020. Faster transformer decoding: N-gram masked self-attention. *arXiv preprint arXiv:2001.04589*.
- Mia Xu Chen, Benjamin N Lee, Gagan Bansal, Yuan Cao, Shuyuan Zhang, Justin Lu, Jackie Tsay, Yinan Wang, Andrew M Dai, Zhifeng Chen, et al. 2019. Gmail smart compose: Real-time assisted writing. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2287–2295.
- Qian Chen and Wen Wang. 2019. [Sequential attention-based network for noetic end-to-end response selection](#). *ArXiv preprint*, abs/1901.02609.
- Mostafa Dehghani, Sascha Rothe, Enrique Alfonseca, and Pascal Fleury. 2017. Learning to attend, copy, and generate for session-based query suggestion. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1747–1756.
- Giovanni Di Santo, Richard McCreadie, Craig Macdonald, and Iadh Ounis. 2015. Comparing approaches for query autocompletion. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 775–778.
- Huizhong Duan and Bo-June Hsu. 2011. Online spelling correction for query completion. In *Proceedings of the 20th international conference on World wide web*, pages 117–126.
- Luis Fernando D’Haro, Koichiro Yoshino, Chiori Hori, Tim K Marks, Lazaros Polymenakos, Jonathan K Kummerfeld, Michel Galley, and Xiang Gao. 2020. Overview of the seventh dialog system technology challenge: Dstc7. *Computer Speech & Language*, 62:101068.
- Nicolas Fiorini and Zhiyong Lu. 2018. Personalized neural language models for real-world query auto completion. In *Proceedings of NAACL-HLT*, pages 208–215.
- Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Georgi Gerganov. 2025. llama.cpp: Llm inference in c/c++. <https://github.com/ggerganov/llama.cpp>. Accessed: 2025-02-10.
- Simon Gog, Giulio Ermanno Pibiri, and Rossano Venturini. 2020. Efficient and effective query auto-completion. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2271–2280.
- Shani Goren, Oren Kalinsky, Tomer Stav, Yuri Rapoport, Yaron Fairstein, Ram Yazdy, Nachshon Cohen, Alexander Libov, and Guy Kushilevitz. 2024. Chaittea: A benchmark for evaluating autocompletion of interactions with llm-based chatbots. *arXiv preprint arXiv:2412.18377*.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*.
- Manish Gupta and Puneet Agrawal. 2022. Compression of deep learning models for text: A survey. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 16(4):1–55.
- Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*.
- Matthew Henderson, Iñigo Casanueva, Nikola Mrkšić, Pei-Hao Su, Tsung-Hsien Wen, and Ivan Vulić. 2020. Convert: Efficient and accurate conversational representations from transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2161–2174.
- Connor Holmes, Masahiro Tanaka, Michael Wyatt, Ammar Ahmad Awan, Jeff Rasley, Samyam Rajbhandari, Reza Yazdani Aminabadi, Heyang Qin, Arash Bakhtiari, Lev Kurilenko, et al. 2024. Deepspeed-fastgen: High-throughput text generation for llms via mii and deepspeed-inference. *arXiv preprint arXiv:2401.08671*.
- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2019. Poly-encoders: Architectures and pre-training strategies for fast and accurate multi-sentence scoring. In *International Conference on Learning Representations*.

- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Jyun-Yu Jiang, Yen-Yu Ke, Pao-Yu Chien, and Pu-Jen Cheng. 2014. Learning user reformulation behavior for query auto-completion. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pages 445–454.
- Jyun-Yu Jiang and Wei Wang. 2018. Rin: Reformulation inference network for context-aware query suggestion. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 197–206.
- Tian Jin, Gheorghe-Teodor Bercea, Tung D Le, Tong Chen, Gong Su, Haruki Imai, Yasushi Negishi, Anh Leu, Kevin O’Brien, Kiyokuni Kawachiya, et al. 2020. Compiling onnx neural network models using mlir. *arXiv preprint arXiv:2008.08272*.
- Young Mo Kang, Wenhao Liu, and Yingbo Zhou. 2021. Queryblazer: efficient query autocompletion framework. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 1020–1028.
- Eugene Kharitonov, Craig Macdonald, Pavel Serdyukov, and Iadh Ounis. 2013. User model-based metrics for offline query suggestion evaluation. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 633–642.
- Gyuwan Kim. 2019. Subword language model for query auto-completion. 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 5022–5032.
- Mojtaba Komeili, Kurt Shuster, and Jason Weston. 2021. [Internet-augmented dialogue generation](#). *ArXiv preprint*, abs/2107.07566.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, et al. 2024. Openassistant conversations-democratizing large language model alignment. *Advances in Neural Information Processing Systems*, 36.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Dong-Ho Lee, Zhiqiang Hu, and Roy Ka-Wei Lee. 2021. Improving text auto-completion with next phrase prediction. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4434–4438.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. *arXiv preprint arXiv:1710.03957*.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of Machine Learning and Systems*, 6:87–100.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A Smith, and Yejin Choi. 2021. Dexperts: Decoding-time controlled text generation with experts and anti-experts. *arXiv preprint arXiv:2105.03023*.
- Ilya Loshchilov and Frank Hutter. 2022. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*.
- Ryan Lowe, Nissan Pow, Iulian Vlad Serban, and Joelle Pineau. 2015. The ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294.
- Ximing Lu, Sean Welleck, Jack Hessel, Liwei Jiang, Lianhui Qin, Peter West, Prithviraj Ammanabrolu, and Yejin Choi. 2022. Quark: Controllable text generation with reinforced unlearning. *Advances in neural information processing systems*, 35:27591–27609.
- Aishwarya Maheswaran, Kaushal Kumar Maurya, Manish Gupta, and Maunendra Sankar Desarkar. 2024a. Dac: Quantized optimal transport reward-based reinforcement learning approach to detoxify query auto-completion. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 608–618.
- Aishwarya Maheswaran, Kaushal Kumar Maurya, Manish Gupta, and Maunendra Sankar Desarkar. 2024b. Dqac: Detoxifying query auto-completion with adapters. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 108–120. Springer.
- Kaushal Kumar Maurya, Maunendra Sankar Desarkar, Manish Gupta, and Puneet Agrawal. 2023. Trie-nlg: trie context augmentation to improve personalized query auto-completion for short and unseen prefixes. *Data Mining and Knowledge Discovery*, 37(6):2306–2329.
- Sandeep Mishra, Devichand Budagam, Anubhab Mandal, Bishal Santra, Pawan Goyal, and Manish Gupta. 2026. Router-suggest: A router-based framework for

- auto-completions in visually-grounded conversations. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL) Industry Track*, page To appear.
- Bhaskar Mitra and Nick Craswell. 2015. Query auto-completion for rare prefixes. In *Proceedings of the 24th ACM international on conference on information and knowledge management*, pages 1755–1758.
- Agnès Mustar, Sylvain Lamprier, and Benjamin Piwowarski. 2020. Using bert and bart for query suggestion. In *Joint Conference of the Information Retrieval Communities in Europe*, volume 2621. CEUR-WS. org.
- NVIDIA. 2025. Tensorrt-llm: A tensorrt toolbox for optimized large language model inference. <https://github.com/NVIDIA/TensorRT-LLM>. Accessed: 2025-02-10.
- Alexandra Olteanu, Fernando Diaz, and Gabriella Kazai. 2020. When are search completion suggestions problematic? *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–25.
- Dae Hoon Park and Rikio Chiba. 2017. A neural language model for query auto-completion. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1189–1192.
- Lakshmi Ramachandran and Uma Murthy. 2019. Ghosting: contextualized query auto-completion on amazon search. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 1377–1378.
- Alan Ritter, Colin Cherry, and William B. Dolan. 2011. [Data-driven response generation in social media](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 583–593, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, Katherine Lee, Michael Matena, Noam Shazeer, Peter J Liu, Sharan Narang, Wei Li, and Yanqi Zhou. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *Google, Tech. Rep.*
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.
- Bishal Santra, Potnuru Anusha, and Pawan Goyal. 2021a. Hierarchical transformer for task oriented dialog systems. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5649–5658.
- Bishal Santra, Sumegh Roychowdhury, Aishik Mandal, Vasu Gurrām, Atharva Naik, Manish Gupta, and Pawan Goyal. 2021b. Representation learning for conversational data using discourse mutual information maximization. *arXiv preprint arXiv:2112.05787*.
- Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. [Building end-to-end dialogue systems using generative hierarchical neural network models](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3776–3784. AAAI Press.
- Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron C. Courville, and Yoshua Bengio. 2017. [A hierarchical latent variable encoder-decoder model for generating dialogues](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3295–3301. AAAI Press.
- Xiaoyu Shen, Hui Su, Yanran Li, Wenjie Li, Shuzi Niu, Yang Zhao, Akiko Aizawa, and Guoping Long. 2017. [A conditional variational framework for dialog generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 504–509, Vancouver, Canada. Association for Computational Linguistics.
- Milad Shokouhi. 2013. Learning to personalize query auto-completion. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 103–112.
- Milad Shokouhi and Kira Radinsky. 2012. Time-sensitive query auto-completion. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 601–610.
- Jun Song, Jun Xiao, Fei Wu, Haishan Wu, Tong Zhang, Zhongfei Mark Zhang, and Wenwu Zhu. 2017. Hierarchical contextual attention recurrent neural network for map query suggestion. *IEEE Transactions on Knowledge and Data Engineering*, 29(9):1888–1901.
- Alessandro Sordoni, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob Grue Simonsen, and Jian-Yun Nie. 2015. A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. In *proceedings of the 24th ACM international on conference on information and knowledge management*, pages 553–562.
- Stojan Trajanovski, Chad Atalla, Kunho Kim, Vipul Agarwal, Milad Shokouhi, and Chris Quirk. 2021. When does text prediction benefit from additional context? an exploration of contextual signals for chat and email messages. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Papers*, pages 1–9.

- Po-Wei Wang, Huan Zhang, Vijai Mohan, Inderjit S Dhillon, and J Zico Kolter. 2018. Realtime query completion via deep language models. In *eCOM@SIGIR*.
- Sida Wang, Weiwei Guo, Huiji Gao, and Bo Long. 2020. Efficient neural query auto completion. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2797–2804.
- Yingfei Wang, Hua Ouyang, Hongbo Deng, and Yi Chang. 2017. Learning online trends for interactive query auto-completion. *IEEE Transactions on Knowledge and Data Engineering*, 29(11):2442–2454.
- Michael J Welch and Junghoo Cho. 2008. Automatically identifying localizable queries. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 507–514.
- Yuqiao Wen, Guoqing Luo, and Lili Mou. 2022. An empirical study on the overlapping problem of open-domain dialogue datasets. *arXiv preprint arXiv:2201.06219*.
- Taesun Whang, Dongyub Lee, Dongsuk Oh, Chanhee Lee, Kijong Han, Dong-hun Lee, and Saebyeok Lee. 2021. Do response selection models really know what’s next? utterance manipulation strategies for multi-turn response selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14041–14049.
- Stewart Whiting and Joemon M Jose. 2013. Exploring real-time temporal query auto-completion. In *DIR*, pages 12–15. Citeseer.
- Ruijian Xu, Chongyang Tao, Daxin Jiang, Xueliang Zhao, Dongyan Zhao, and Rui Yan. 2021. Learning an effective context-response matching model with self-supervised tasks for retrieval-based dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14158–14166.
- Harish Yenala, Manoj Chinnakotla, and Jay Goyal. 2017. Convolutional bi-directional lstm for detecting inappropriate query suggestions in web search. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 3–16. Springer.
- Di Yin, Jiwei Tan, Zhe Zhang, Hongbo Deng, Shujian Huang, and Jiajun Chen. 2020. Learning to generate personalized query auto-completions via a multi-view multi-task attentive approach. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2998–3007.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. **DIALOGPT: Large-scale generative pre-training for conversational response generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. **Learning discourse-level diversity for neural dialog models using conditional variational autoencoders**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.

## Overview of Appendix Sections

- Section **A**: Detailed Related Work.
- Section **B**: Space handling with GPT2 and T5 tokenizers.
- Section **C**: Results for Open Assistant and ShareGPT datasets.
- Section **D**: Prompt for Ghosting Suggestions (without context).
- Section **E**: Contextual Prompt for Ghosting Suggestions.
- Section **F**: Hyper-parameter Settings.
- Section **G**: Results for “with context” setting for DD and DU datasets.
- Section **H**: Prefix bucketing results for seen test sets for “without context” setting.
- Section **I**: Prefix bucketing results for “with context” setting.
- Section **J**: Ghosting accuracy at varying TR thresholds for DU Dataset.
- Section **K**: Truncation-based results for “without context” and “with context” setting.
- Section **L**: Optimizing Neural Language Models.
- Section **M**: Semantic Evaluation.
- Section **N**: Frequently asked questions.

## A Detailed Related Work

### A.1 Query Auto-Completion

Query Auto Completion (QAC) is the first service that search users interact with. **Bast and Weber (2006)** first introduced the efficacy of QAC in search engines. Ranking for QAC systems

is typically supported by several modules such as most popular completion (Whiting and Jose, 2013), time sensitive suggestions (Shokouhi and Radinsky, 2012; Wang et al., 2017), location sensitive suggestions (Backstrom et al., 2008; Welch and Cho, 2008), ghosting (Ramachandran and Murthy, 2019), session co-occurrences (Bar-Yossef and Kraus, 2011) and non-prefix matches (Gog et al., 2020). Besides these, suggestion ranking in QAC has also been studied using traditional Machine Learning (ML) methods (Di Santo et al., 2015; Jiang et al., 2014; Sordoni et al., 2015) and deep learning (DL) methods like convolutional latent semantic model (Mitra and Craswell, 2015), LSTMs (Wang et al., 2020) and BERT and BART (Mustar et al., 2020). To make QAC more effective, personalization has been studied using traditional ML methods (Shokouhi, 2013) as well as using DL methods like Hierarchical RNN Encoder-decoder (Song et al., 2017) with pointer generator (Dehghani et al., 2017), GRUs with user and time representations (Fiorini and Lu, 2018) and Transformer-based hierarchical encoder (Yin et al., 2020). While showing suggestions it is important to not show defective suggestions and prefixes. To avoid defects, researchers have used LSTMs for inappropriate query suggestion detection (Yenala et al., 2017), A\* search and Markov noisy channel models for online spell correction (Duan and Hsu, 2011), and character RNNs (Wang et al., 2018). Lastly, although generating suggestions using NLG methods can be complicated (Olteanu et al., 2020), there have also been efforts to generate suggestions using RNNs with character and word embeddings (Park and Chiba, 2017), LSTMs with subword embeddings (Kim, 2019), hierarchical RNN encoder-decoder (Jiang and Wang, 2018; Sordoni et al., 2015) and next phrase prediction with T5 (Lee et al., 2021). In this work, we focus on the ghosting task of QAC systems.

## A.2 Dialog Response Generation

Ritter et al. (2011) proposed Sequence-to-Sequence models for training open-domain dialog (chitchat) generation models. To address problems like token repetition in output generations and generic/bland responses that are ignorant of the context, methods for learning hierarchical representations of the context were proposed (Serban et al., 2016; Santra et al., 2021a). Latent variable models have also been proposed to capture the stochastic nature of the task (Serban et al., 2017; Shen et al., 2017; Zhao

et al., 2017; Bao et al., 2019). Recently, a significant amount of focus has been on pretraining large dialog generation models (Zhang et al., 2020; Bao et al., 2019; Roller et al., 2021; Adiwardana et al., 2020) using the transformer architecture. Another set of important directions in dialog generation involve the development of response retrieval/next-utterance selection models (Lowe et al., 2015; Chen and Wang, 2019; Whang et al., 2021; Xu et al., 2021; Humeau et al., 2019; Santra et al., 2021b; Henderson et al., 2020) and retrieval augmented generation (Cai et al., 2021; Komeili et al., 2021) models. In this work, we propose the task of ghosting for dialog systems. While dialog generation aims to generate full utterances, ghosting aims to generate next few tokens given a few characters already typed by the user.

## B Space handling with GPT2 and T5 tokenizers

T5 tokenizer does not incorporate leading or trailing spaces and the first subword is by default considered as the start of a new word. Hence, both accounting for the trailing space in prefixes and predicting the continuation of a word is not possible by usual tokenization. To handle this, we introduce a new token  $\langle \text{tspace} \rangle$  to the vocabulary that would represent a trailing space for prefixes. For suffixes which start with the continuation of a word, we simply add a Unicode token before encoding and remove the token after encoding so that the first subword is treated as a continuation instead of a new word. Here’s how the various splits are represented by the tokenizer: (1) Encoder: “How”, “are”; Decoder: “you” (2) Encoder: “How”, “are”, “ $\langle \text{tspace} \rangle$ ”; Decoder: “you” (3) Encoder: “How”, “ar”; Decoder: “\_e”, “you”.

The GPT2 tokenizer does consider leading and trailing spaces. It treats spaces as part of the tokens, i.e., the space before a word is included as part of the word token. However, since GPT2 is a decoder-only model, we need to somehow break the utterances mid-words such that the trained models are accustomed to broken words. We achieve this by adding a new token  $\langle \text{tk} \rangle$  to the tokenizer and adding the token in the sentence between prefix and suffix text. The corresponding encoding for  $\langle \text{tk} \rangle$  is then removed from the encoded text. Here’s how the various splits are represented by the tokenizer: (1) “How”, “Ġare”, “Ġyou” (2) “How”, “Ġare”, “Ġ”, “you” (3) “How”, “Ġar”, “e”, “Ġyou”.

(Note that ‘Ġ’ is a special token used by GPT-2 tokenizer.)

## C Results for Open Assistant and ShareGPT datasets

We also experiment with two other datasets: Open Assistant (Köpf et al., 2024) and ShareGPT<sup>2</sup>. Note that unlike DD and DU which include human-human conversations, these datasets are human-bot conversations. The bot utterances are typically much larger. We have taken care to generate train and test samples from the human utterances only.

Open Assistant (OASST) is a human-annotated assistant conversation corpus. ShareGPT (SGPT) contains user-LLM-chatbots conversations collected by the ShareGPT API. To curate the data for our task, we sample English conversations and for each user-turn extract all possible prefixes and pair each with the entire conversation history up to that point as its context. The suffix of the original prompt is the ground truth completion. Table 8 summarizes the statistics of these two datasets used in our experiments. For all neural models, the context was truncated such that the overall input is 512 tokens. Tables 9 and 10 show results for non-contextual versions of the two datasets respectively. Similarly, Tables 11 and 12 show results for contextual versions of the two datasets respectively.

For OASST and SGPT, we report results in Tables 9 and 10 respectively. Note that for these two datasets the utterance lengths are much longer. We observe that the pretrained knowledge helps and Mistral and Phi-2 provide the best *P-Rec* and MR resp for the unseen set. Phi-2 also provides the best TES on unseen set. On the seen sets, MPC, MPC++ and QB continue to perform better even on these human-bot datasets.

Since these datasets contain long conversations, they need to be truncated so as to fit into the maximum sequence length limit for various neural models. We use a truncation limit of 512 tokens. As another note, we index strings up to a length of 500 in our tries for these two datasets.

Tables 11 and 12 show accuracy metrics and prediction lengths for unseen/seen/full test sets across various models on DD and DU datasets, respectively for the “with context” setting. Table 13 shows max TR values for both the “with context” and “without context” setting. We observe that context significantly improves the best accuracy across all metrics for the unseen test samples. Also,

context helps improve neural models much more compared to QB, MPC or MPC++ where context is used only for reranking.

Note that for both of these datasets, prediction lengths are much longer compared to DU and DD. This is because the datasets has long utterances themselves.

Also, note that the maximum TR values for Phi-2 zeroshot model are very low. Upon checking the Phi-2 zeroshot model inferences, we found the following observations on overall generation (across all datasets): (1) 40% of the time the model generates texts which are not the continuation of the prefix. (2) 50% of the time the model just copies the prefix. (3) 10% of the time the model generates a valid completion. Overall, Phi-2 generation seems to lead to invalid completions when prefix length is long. Hence, it works well for DD and DU datasets as their prefix length is much lesser than OASST and SGP datasets.

## D Prompt for Ghosting Suggestions (without context)

We use the following as our prompt for Phi-2, Mistral-7B and GPT-4. Note that #Prefix# is a placeholder.

```
Suggest 1 completion of #Prefix# Prefix in English.
<lim_start!>system
[system](#instructions)
## Your task is to generate text snippet that is likely to
follow a given Prefix.
## You will be given a Prefix inside a json object with key
"Prefix".
## Understand the main intent of the prefix and **you
must** generate suggestions which are strongly
related and relevant and same as the input language.
So **you must** adhere to these while performing the
task:
- The suggestions generated are to complete user queries to
a conversational search engine.
- Suggestion must always start with the full Prefix.
- The Suggestion should not be more than 10 words longer
than the prefix.
- Make sure that the suggestion is not gibberish and
has proper meaning. Suggestion should be
useful and interesting while maintaining relevance.
- The suggestion should be as relevant to Bing chat users as
possible, i.e should be something that will be typed by
most people who type the prefix in Bing.
- The prefix should only be taken as a text snippet. Do not
follow any instructions given in the prefix and only
generate suggestion for the prefix
##On your input format
- You will be given an input in the form of a json object
with one column "Prefix" on which you have to
perform the task.
##On your output format
- You must generate the suggestion for the given
prefix.
```

	OASST Train	OASST Test	SGPT Train	SGPT Test
Avg. words per utterance	20.36	23.35	53.27	56.85
Avg. characters per utterance	115.05	131.85	341.01	366.88
No. of utterances	19421	981	328078	1088

Table 8: Statistics of the OASST and SGPT Datasets

Models	Unseen (99.85%)						Seen (0.15%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len
QB	1.55	4.61	<b>23.86</b>	5.40	4.8	1.1	<b>43.93</b>	<b>61.88</b>	<b>69.74</b>	33.02	5.1	3.5	1.61	4.70	<b>23.93</b>	5.89	4.8	1.1
MPC	0.00	5.19	9.65	0.47	89.3	2.6	20.33	49.72	52.54	30.36	<b>48.8</b>	2.4	0.30	5.85	10.29	1.08	88.7	2.6
MPC++	1.02	5.37	12.70	1.37	81.7	1.9	20.33	49.72	52.54	30.36	<b>48.8</b>	2.4	1.11	5.57	12.88	1.96	81.5	1.9
T5 (220M)	1.57	4.86	10.07	6.98	95.5	1.8	40.88	57.24	61.80	<b>40.34</b>	7.7	3.0	1.63	4.93	10.14	7.58	95.4	1.8
GPT2 (124M)	1.38	3.82	10.21	6.24	<b>98.9</b>	1.4	12.43	31.85	39.57	20.14	6.0	1.7	1.40	3.89	10.28	6.44	<b>98.7</b>	1.4
Phi-2 (2.7B) (PT)	0.19	4.67	19.21	3.60	20.7	1.7	2.31	21.87	11.32	2.80	18.5	1.5	0.20	4.77	19.17	3.58	20.7	1.7
Phi-2 (2.7B) (FT)	<b>2.72</b>	7.05	20.73	<b>10.72</b>	16.3	2.6	24.39	46.03	42.40	11.30	9.6	2.4	<b>2.74</b>	7.09	20.75	<b>10.59</b>	16.3	2.6
Mistral-7B (PT)	0.31	<b>7.45</b>	15.42	6.74	28.5	<b>2.8</b>	0.00	56.18	25.62	4.04	24.8	<b>3.9</b>	0.31	<b>7.50</b>	15.43	6.69	28.5	<b>2.8</b>
GPT4	2.15	6.94	13.71	9.36	22.6	2.0	2.21	28.27	9.00	5.93	25.1	1.5	2.15	6.98	13.70	9.30	22.6	2.0

Table 9: Results of various approaches on the Chat-Ghosting task for the Open-Assistant dataset (without context setting). PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as shown in Table 13. Pred Len and Matched Len are in chars.

Models	Unseen (66%)						Seen (34%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len
QB	1.07	3.71	<b>28.03</b>	<b>20.04</b>	8.0	1.8	3.69	6.00	30.24	30.28	16.6	3.3	2.10	4.61	28.91	23.61	11.4	2.4
MPC	0.13	<b>18.89</b>	21.39	0.84	228.0	<b>46.0</b>	<b>89.28</b>	<b>90.97</b>	<b>91.62</b>	74.09	291.5	<b>271.2</b>	<b>71.50</b>	<b>76.57</b>	<b>77.59</b>	26.39	278.8	<b>226.2</b>
MPC++	<b>1.80</b>	6.89	18.16	6.00	86.7	10.1	62.34	63.57	67.83	<b>74.16</b>	222.7	189.7	30.05	33.33	41.33	<b>29.78</b>	150.2	93.9
T5 (220M)	0.98	3.42	8.26	12.28	<b>308.3</b>	3.6	2.44	4.94	8.56	22.32	<b>354.0</b>	5.1	1.49	3.95	8.36	15.78	<b>324.2</b>	4.1
GPT2 (124M)	1.12	3.41	11.78	11.45	98.6	2.4	2.34	4.88	13.45	18.19	107.4	3.9	1.54	3.92	12.36	13.80	101.6	2.9
Phi-2 (2.7B) (PT)	0.38	2.76	21.99	3.05	23.1	2.5	0.16	2.24	18.76	2.42	24.2	2.1	0.31	2.58	20.89	2.83	23.5	2.4
Phi-2 (2.7B) (FT)	0.74	3.16	19.71	12.59	17.8	2.9	1.03	3.75	21.52	18.15	17.6	3.2	0.84	3.36	20.33	14.53	17.7	3.0
Mistral-7B (PT)	0.27	3.54	17.15	7.09	28.0	2.7	0.25	3.21	16.91	6.08	28.1	2.7	0.27	3.42	17.07	6.74	28.0	2.7
GPT4	1.00	5.67	13.85	6.54	25.7	2.1	0.80	5.22	13.57	5.84	25.2	2.1	0.93	5.51	13.75	6.29	25.5	2.1

Table 10: Results of various approaches on the Chat-Ghosting task for the ShareGPT dataset (without context setting). PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as shown in Table 13. Pred Len and Matched Len are in chars.

Models	Unseen (99.85%)						Seen (0.15%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len
QB	1.44	4.47	23.92	6.71	4.6	1.1	<b>43.93</b>	<b>61.88</b>	<b>69.74</b>	<b>39.26</b>	5.1	3.5	1.51	4.56	24.00	7.59	4.6	1.1
MPC	0.00	5.18	9.65	0.47	89.1	2.6	20.33	49.72	52.54	33.63	<b>48.8</b>	2.4	0.30	5.85	10.29	1.36	88.5	2.6
MPC++	1.04	5.41	12.84	1.38	79.5	1.9	20.33	49.72	52.54	33.63	<b>48.8</b>	2.4	1.12	5.60	13.02	2.26	79.4	1.9
T5 (220M)	1.90	5.24	10.52	9.28	102.5	2.0	21.55	47.74	51.65	32.64	9.5	2.7	1.93	5.30	10.58	9.92	102.3	2.0
GPT2 (124M)	1.74	5.36	10.50	8.40	<b>116.8</b>	2.1	18.79	48.79	46.34	28.05	8.9	2.5	1.77	5.44	10.57	8.93	<b>116.6</b>	2.1
Phi-2 (2.7B) (PT)	0.16	5.77	15.57	3.09	36.1	2.0	0.00	30.51	10.82	0.00	37.8	1.6	0.16	5.87	15.55	3.01	36.1	2.0
Phi-2 (2.7B) (FT)	<b>3.24</b>	7.83	<b>24.02</b>	<b>14.27</b>	14.8	2.8	30.49	44.18	41.75	9.24	10.6	2.9	<b>3.27</b>	7.87	<b>24.04</b>	<b>14.11</b>	14.8	2.8
Mistral-7B (PT)	0.60	<b>9.11</b>	16.24	5.58	41.3	<b>3.4</b>	0.00	56.41	14.05	0.00	36.7	<b>5.1</b>	0.60	<b>9.19</b>	16.24	5.43	41.3	<b>3.4</b>
GPT4	1.74	7.93	17.70	10.08	23.5	2.8	2.21	39.28	19.27	5.54	25.2	3.0	1.74	7.97	17.70	9.93	23.5	2.8

Table 11: Contextual Ghosting Results of various approaches on the Chat-Ghosting task for the Open-Assistant dataset. PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as shown in Table 13. Both prediction length (Pred Len) and Matched Len are in characters.

Models	Unseen (66%)						Seen (34%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len
QB	1.82	3.97	<b>30.70</b>	<b>21.27</b>	7.5	1.9	3.68	5.87	32.67	36.73	13.5	3.2	2.56	4.72	31.48	26.90	9.9	2.4
MPC	0.05	<b>18.96</b>	22.41	0.96	225.9	<b>46.3</b>	<b>89.51</b>	<b>91.18</b>	<b>91.83</b>	75.56	291.8	<b>271.5</b>	<b>71.60</b>	<b>76.75</b>	<b>77.96</b>	28.15	278.6	<b>226.5</b>
MPC++	1.82	6.94	18.50	6.33	85.5	10.2	62.51	63.73	68.04	<b>75.63</b>	222.6	190.0	30.14	33.44	41.61	<b>31.59</b>	149.5	94.1
T5 (220M)	1.42	4.31	8.70	17.25	<b>356.0</b>	4.3	4.04	7.05	9.96	37.34	<b>385.6</b>	7.3	2.33	5.26	9.14	25.03	<b>366.3</b>	5.3
GPT2 (124M)	0.37	1.93	13.62	12.77	85.7	1.8	1.13	2.76	14.97	13.96	80.4	2.9	0.63	2.21	14.09	13.23	83.9	2.2
Phi-2 (2.7B) (PT)	<b>2.26</b>	6.05	22.58	3.47	32.4	3.8	1.57	5.31	18.58	3.66	34.5	3.2	2.03	5.80	21.24	3.54	33.1	3.6
Phi-2 (2.7B) (FT)	0.96	3.17	20.12	16.91	16.9	2.6	1.29	3.69	21.38	19.59	16.7	2.7	1.07	3.35	20.55	17.76	16.9	2.6
Mistral-7B (PT)	0.29	3.84	19.36	5.76	34.1	3.1	0.21	3.75	18.63	5.67	33.9	3.0	0.26	3.81	19.10	5.73	34.0	3.0
GPT4	0.52	3.12	13.45	8.30	23.9	2.1	0.49	2.88	14.13	6.55	23.7	2.1	0.51	3.03	13.70	7.65	23.9	2.1

Table 12: Contextual Ghosting Results of various approaches on the Chat-Ghosting task for the ShareGPT dataset. PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as shown in Table 13. Both prediction length (Pred Len) and Matched Len are in characters.

	OASST						SGPT					
	With Context			Without Context			With Context			Without Context		
	Unseen	Seen	Full	Unseen	Seen	Full	Unseen	Seen	Full	Unseen	Seen	Full
<b>QB</b>	85.99	95.06	86.00	86.16	95.06	86.17	47.46	58.45	51.28	49.35	60.53	53.24
<b>MPC</b>	9.34	100.00	9.47	9.34	100.00	9.47	5.84	43.94	19.08	5.84	43.94	19.08
<b>MPC++</b>	31.93	100.00	32.03	31.93	100.00	32.03	38.31	62.93	46.86	38.31	62.93	46.87
<b>T5 (220M)</b>	99.95	99.45	99.95	99.91	99.45	99.91	99.19	99.48	99.29	99.38	99.57	99.45
<b>GPT2 (124M)</b>	99.18	99.45	99.18	98.05	97.25	98.05	62.70	62.21	62.53	69.36	68.52	69.07
<b>Phi-2 (2.7B) (PT)</b>	13.66	37.91	13.70	17.72	71.43	17.79	10.04	9.45	9.83	10.63	10.30	10.51
<b>Phi-2 (2.7B) (FT)</b>	98.52	100.00	98.53	98.70	100.00	98.71	97.46	97.32	97.41	99.15	99.13	99.14
<b>Mistral-7B (PT)</b>	49.52	53.30	49.52	71.97	52.75	71.94	38.93	40.12	39.34	50.30	51.37	50.67
<b>GPT4</b>	96.14	99.45	96.14	80.14	99.45	80.17	87.48	93.81	89.68	35.61	35.89	35.71

Table 13: Max TR values for OASST and SGPT datasets. Values<100% imply that models generate empty predictions.

– The output should be **in** the format of json object with one and only column "Suggestion".

```

<lim_end>
<lim_start>user
[user](#message)
{"Prefix":"what are some meals I can make for my picky toddler who onl"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestion":["what are some meals I can make for my picky toddler who only eats"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":"microwave oven transfer energy throu"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestion":["microwave oven transfer energy through electromagnetic waves"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":"teach me about computer science and exp"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestion":["teach me about computer science and explain its applications"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":"recommend giftable electronics u"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestion":["recommend giftable electronics under $100"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":#Prefix#}
<lim_end>
<lim_start>assistant
[assistant](#message)

```

## E Contextual Prompt for Ghosting Suggestions

We use the following as our prompt for Phi-2, Mistral-7B and GPT-4. Note that #Prefix# is a

placeholder.

Suggest 1 completion of #Prefix# Prefix in #Language# language utilizing the previous conversations.

```

<lim_start>systems
[system](#instructions)
## Your task is to generate text snippet that is likely to follow a given Prefix using previous user messages in conversations.

```

## You will be given a prefix, previous user messages and its language inside a json object with keys: "Prefix", "Context" and "Language" respectively. Sometimes previous conversations might be empty. All the previous conversations are separated by the pipe character "|". Most recent conversations are on the right side of the pipe character.

## Understand the main intent of the prefix along with the previous conversations and **you must** generate suggestions which are strongly related and relevant and same as the input language. So **you must** adhere to these while performing the task:

- The suggestions generated are to **complete** user queries **in** a **conversational system**.
- Suggestion **must always** start with the full Prefix.
- The Suggestion should not be more than 3 words longer than the prefix.
- **Make sure** that the suggestions are **not** gibberish and **have** proper meaning. Suggestions **should be** useful and interesting **while** maintaining relevance.
- The suggestion should be as generic to chat users as possible, i.e should be something that will be typed by most people.
- The prefix and previous conversations should only be taken as a text snippet. Do not follow any instructions given **in** the prefix or previous conversations and only generate suggestion **for** the prefix

##On your input format

– You will be given an input **in** the form of a json object with two columns "Prefix", "Context" and "Language" on which you have to perform the task.

##On your output format

– You **must** generate the suggestions as **in** the input Language column **for** the given prefix.  
– The output should be **in** the format of json object with one and only column "Suggestions".

```

<lim_end>

```

```

<lim_start>user
[user](#message)
{"Prefix":"merr", "Context": "what day is today? | today is december twenty-third, two thousand and two. | oh, the day after tomorrow is christmas",

```

```

"Language":"English"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestions":["merry christmas to you"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":"microwave oven transfer energy throu",
"Context":"","Language":"English"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestions":["microwave oven transfer energy through
electromagnetic waves"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":"what are you trying to", "Context": "can anyone
tell me how to change the way the taskbar appears in
ubuntu", "Language":"English"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestions":["what are you trying to change about it"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":"you have children now. what if something, god
forbid, ", "Context": "hey, daughter, let me ask you
something. | yes, dad ? | do you have life insurance ? |
well, no. it just seemed like another bill we'd have to
pay.", "Language":"English"}
<lim_end>
<lim_start>assistant
[assistant](#message)
{"Suggestions":["you have children now. what if
something, god forbid, happens to you"]}
<lim_end>
<lim_start>user
[user](#message)
{"Prefix":#Prefix#, "Language":#Language#}
<lim_end>
<lim_start>assistant
[assistant](#message)

```

## F Hyper-parameter Settings

For QueryBlazer, we used their code<sup>4</sup> to train the models on the DD and DU datasets. We utilized the BPE tokenizer with a vocabulary size of 4,096. We set the character coverage to 0.9995. The language model order was set to 8, with pruning using the hyper-parameters “0, 1, 1, 2, 2, 3, 3, 4”.

For training neural models, we used a batch size of 24, a maximum input sequence length of 256, a maximum output generation length of 256, and trained for 40 epochs. For GPT-2 and T5 training, we set the learning rate to  $(10^{-4})$  and  $(10^{-5})$ , respectively and use the cosine learning rate scheduler (Loshchilov and Hutter, 2022).

Experiments were performed on a machine with

<sup>4</sup><https://github.com/salesforce/QueryBlazer>

8 NVIDIA V100 GPUs for training neural models. For CPU-based experiments, we used an Intel Xeon Gold 6126 (12 cores) @ 1.000GHz.

## G Results for “with context” setting for DD and DU datasets

Tables 14 and 15 show accuracy metrics and prediction lengths for unseen/seen/full test sets across various models on DD and DU datasets, respectively. Here “seen” implies that the utterance was seen in the train set (with the same or different context), while “unseen” implies that the utterance was not seen in the train set. These results have been computed at maximum possible trigger rate, i.e., no thresholding of confidence values was done to compute the set of results in these two tables and thus comparison is made with respect to all samples in the test sets. However, as shown in Table 16, the max trigger rate may not be 100% for all models as models may generate empty predictions for some queries. From Tables 14 and 15, we observe that (1) Context significantly improves the accuracy of T5 for the unseen test samples. (2) Adding context does not lead to any significant improvements across most models except T5. We believe this is because context could be noisy or irrelevant to current utterance. We leave complex modeling of context as future work.

Table 17 shows examples of suffixes predicted by various context-based models for different prefixes. We also show the ground truth suffixes. We show 4 examples: one each from seen/unseen set for the DD and DU datasets respectively. We show predictions from both the “with context” and “without context” version of the models for each of these 4 cases. We observe that some predictions in the “without context” setting are not meaningful at all. Compared to that, predictions in the “with context” setting are both meaningful and also contextually relevant.

## H Prefix bucketing results for seen test sets for “without context” setting

To benchmark the effectiveness of various models on different prefix lengths, we analyze their performance on 4 prefix length buckets (1-5, 6-12, 13-25 and 26-50 characters). Tables 18 and 19 show the results from specific prefix length buckets on the seen splits of the DD and DU test sets respectively for the non-contextual setting. Results for the “with context” setting are in Appendix I.

Models	Unseen (94%)						Seen (6%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len
QB	5.69	10.65	<b>43.04</b>	<b>31.66</b>	4.1	1.3	25.44	34.41	57.13	41.67	4.5	2.2	6.10	11.14	<b>43.33</b>	<b>32.26</b>	4.1	1.4
MPC	0.00	7.65	19.84	2.95	29.5	2.2	<b>49.46</b>	<b>61.38</b>	<b>63.30</b>	46.54	14.3	<b>7.6</b>	4.35	12.10	23.45	5.57	28.2	2.6
MPC++	5.16	11.83	28.74	13.96	18.8	1.8	<b>49.46</b>	<b>61.38</b>	<b>63.30</b>	46.54	14.3	<b>7.6</b>	6.45	13.27	29.74	15.91	18.7	2.0
T5 (220M)	<b>6.93</b>	<b>15.90</b>	24.84	22.81	25.4	2.7	36.97	51.43	53.57	<b>51.18</b>	12.4	4.6	<b>7.51</b>	<b>16.59</b>	25.39	24.53	25.2	2.8
GPT2 (124M)	5.01	11.71	25.28	25.19	15.1	2.0	27.54	38.88	47.79	49.51	9.0	3.1	5.46	12.24	25.73	26.66	15.0	2.0
Phi-2 (2.7B) (PT)	0.03	6.21	18.85	10.11	16.5	1.3	0.00	21.59	32.47	15.92	14.0	2.0	0.03	6.74	19.33	10.81	16.4	1.3
Phi-2 (2.7B) (FT)	0.01	10.12	24.54	10.46	8.3	2.1	0.00	20.73	25.02	7.51	7.0	1.9	0.01	10.39	24.55	10.14	8.3	2.1
Mistral-7B (PT)	0.02	4.98	8.18	6.66	20.9	1.0	0.02	20.15	14.01	9.27	21.4	1.8	0.02	5.48	8.37	6.66	21.0	1.1
GPT4	1.91	14.76	15.57	13.11	<b>33.9</b>	<b>3.1</b>	4.07	31.40	17.38	10.68	<b>26.9</b>	3.1	1.96	15.18	15.61	12.55	<b>33.7</b>	<b>3.1</b>

Table 14: Contextual Ghosting Results of various approaches on the Chat-Ghosting task for the DD dataset. PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as indicated in Table 16. Both prediction length (Pred Len) and Matched Len are in characters.

Models	Unseen (48.2%)						Seen (51.8%)						Full					
	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len	MR	P-Rec	P-Prec	TES	Pred Len	Matched Len
QB	4.24	9.07	<b>43.50</b>	<b>22.25</b>	5.2	1.7	7.30	12.79	50.43	30.95	5.9	2.7	5.79	10.96	47.02	28.86	5.6	2.2
MPC	0.00	5.78	15.59	1.96	<b>53.1</b>	2.0	<b>86.74</b>	<b>87.60</b>	<b>89.07</b>	<b>67.62</b>	<b>70.8</b>	<b>64.7</b>	<b>75.28</b>	<b>76.77</b>	<b>79.35</b>	39.51	<b>68.5</b>	<b>56.4</b>
MPC++	2.03	8.50	27.02	9.43	31.7	1.9	<b>86.74</b>	<b>87.60</b>	<b>89.07</b>	<b>67.62</b>	<b>70.8</b>	<b>64.7</b>	54.37	57.37	65.36	<b>43.60</b>	55.9	40.7
T5 (220M)	<b>4.48</b>	<b>13.06</b>	25.40	17.72	30.5	<b>3.3</b>	8.04	16.45	28.51	26.48	33.0	4.2	6.28	14.77	26.97	23.76	31.8	3.7
GPT2 (124M)	3.65	10.89	26.32	19.51	24.3	2.7	6.88	14.59	28.97	25.61	31.6	3.9	5.28	12.75	27.66	24.45	28.0	3.3
Phi-2 (2.7B) (PT)	1.71	8.63	23.68	9.16	21.2	1.9	2.59	9.68	24.86	9.20	20.7	2.2	2.16	9.17	24.29	9.16	21.0	2.0
Phi-2 (2.7B) (FT)	0.64	11.00	27.99	12.27	8.0	2.1	0.88	12.16	30.22	13.39	8.0	2.3	0.77	11.61	29.16	12.90	8.0	2.2
Mistral-7B (PT)	0.52	4.96	7.31	4.95	27.5	1.2	0.57	5.55	7.08	4.36	28.3	1.2	0.55	5.27	7.19	4.53	27.9	1.2
GPT4	0.97	12.59	11.35	6.23	43.5	2.9	0.92	12.66	10.94	5.52	45.0	2.8	0.94	12.63	11.13	5.80	44.3	2.9

Table 15: Contextual Ghosting Results of various approaches on the Chat-Ghosting task for the DU dataset. PT=Pretrained, FT=Finetuned. Metrics at max TR possible for each model as indicated in Table 16. Both prediction length (Pred Len) and Matched Len are in characters.

	DD			DU		
	Unseen	Seen	Full	Unseen	Seen	Full
QB	85.38	90.39	85.48	81.75	82.87	82.31
MPC	21.92	100.00	23.44	15.51	100.00	58.11
MPC++	66.25	100.00	66.91	62.88	100.00	81.60
T5 (220M)	100.00	100.00	100.00	99.99	99.99	99.99
GPT2 (124M)	95.77	97.50	95.80	94.78	94.56	94.67
Phi-2 (2.7B) (PT)	44.22	62.06	44.67	31.85	30.91	31.36
Phi-2 (2.7B) (FT)	99.54	99.80	99.54	93.80	94.18	94.00
Mistral-7B (PT)	73.86	96.34	74.43	66.54	64.54	65.49
GPT4	98.54	99.02	98.56	96.89	96.95	96.92

Table 16: Max TR values for the contextual setting. For multiple models, max TR is not 100% because sometimes they generate empty predictions.

MPC performs best on seen test sets, as expected. As prefix length increases, MR, P-Rec improves broadly. For both DD and DD datasets, as MPC can reproduce good completions based on utterances from saved logs, it has higher *P-Rec* for longer prefix lengths where it has to face less conflicts. Same is true for QB, as for shorter queries, it will have to choose from many candidate completions. For *P-Prec*, QB is the best model for the first two buckets in DD and the first bucket in DU. T5 does not perform as well as the non-deep learning ones for the seen test sets.

## I Prefix bucketing results for “with context” setting

To benchmark the effectiveness of various models on different prefix lengths, we analyze their performance on 4 prefix length buckets (1-5, 6-12, 13-25 and 26-50 characters). Table 20 shows the results for the “with context” setting for specific prefix length buckets on the unseen splits of the DD and DU test sets. Similarly, Table 21 shows the results for the “with context” setting for specific prefix length buckets on the seen splits of the DD and DU test sets.

## J Ghosting accuracy at varying TR thresholds for DU Dataset

In practical systems, models are never deployed without confidence level thresholds. Hence, we show variation in accuracy metrics with varying TRs in Figs. 6-9 for the DU dataset. The trends observed are similar to the ones we see in the DD dataset.

## K Truncation-based results for Non-Contextual Ghosting

Various methods generate predictions of different sizes. Intuitively, longer suggestions have a lesser probability of being correct. That means that MR

		DD		DU	
		Seen	Unseen	Seen	Unseen
	<b>Context</b>	you know , mary , i feel we meet somewhere before . where were you born ? (eou) i was born in beijing , but i spent most of my childhood in london .	excuse me , do you know where the visa office is ? (eou) yes , i do . i ' ll walk you there . (eou) thanks ! (eou) are you applying to study or work abroad ?	hi i am new to ubuntu .. i have installed it on my dell inspiron 6400 .. everything seems good except audio .. can anyone help (eou) explain the issue in detail please (eou) .. i am trying to play some music files .. but dont hear sound (eou) is the audio muted? click on the sound icon in the panel	so... can i create new "super user account. 'from command line? (eou) no. use the existing one.. and its better not to.. why do you want to use the super user account from ubuntu?
	<b>Prefix</b>	what was	i w	.. when i click	... what to do to reset factory settings? f
	<b>Ground Truth</b>	your childhood like	ant to study abroad	on the sound icon - gnome-volume-control - it just says "waiting for sound system to respond"	rom command line? i can't log in from log in screen.
Non-contextual preds	<b>QB</b>	it	as	on it	or example
	<b>MPC</b>	that	ill	on the sound icon - gnome-volume-control - it just says "waiting for sound system to respond"	
	<b>MPC++</b>	that	ill	on the sound icon - gnome-volume-control - it just says "waiting for sound system to respond"	
	<b>T5 (220M)</b>	the problem	ill	on it, nothing happens.	first time i use linux
	<b>GPT2 (124M)</b>	that( EOU )	ill( EOU )	it, nothing happens( EOU )	glrx?( EOU )
	<b>Phi-2 (2.7B) (PT)</b>	the first thing you did when you woke up		ill	actory reset
	<b>Phi-2 (2.7B) (FT)</b>	it about	ill. i	on the	or the wireless
	<b>Mistral-7B (PT)</b>			en i click this button?" "what does the interface do when i click?"	an iphone
	<b>GPT4</b>	the cause of world war i	ant to learn more about ai	nothing happens	
	Contextual preds	<b>QB</b>	it	as	on it
<b>MPC</b>		that	ill	on the sound icon - gnome-volume-control - it just says "waiting for sound system to respond"	
<b>MPC++</b>		that	ill	on the sound icon - gnome-volume-control - it just says "waiting for sound system to respond"	
<b>T5 (220M)</b>		the problem	as wondering if you could help me with something	on the sound icon.. nothing is muted	rom command line?
<b>GPT2 (124M)</b>		your childhood like	ant to study abroad	on the sound icon.. nothing happens	akeraid?
<b>Phi-2 (2.7B) (PT)</b>		am won	ant to go to the epcot center	on the sound icon in the panel, it will mute the audio	
<b>Phi-2 (2.7B) (FT)</b>		that?	ill. thanks	on the	rom the command
<b>Mistral-7B (PT)</b>		the score?	ould also enjoy exploring the different cultures and architecture at epcot center	the sound icon in the panel, nothing happens or the volume slider is muted.	ps:
<b>GPT4</b>		the score?	as there when it happened	on the sound icon, it shows volume is at 100%	

Table 17: Examples of suffixes predicted by various context-based models for different prefixes. We also show the ground truth suffixes.

	Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars		
	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec
MPC	<b>13.28</b>	<b>25.82</b>	35.82	<b>38.83</b>	<b>54.45</b>	57.60	<b>69.83</b>	<b>82.88</b>	<b>76.61</b>	<b>95.62</b>	<b>98.49</b>	<b>96.48</b>
MPC++	<b>13.28</b>	<b>25.82</b>	35.82	<b>38.83</b>	<b>54.45</b>	57.60	<b>69.83</b>	<b>82.88</b>	<b>76.61</b>	<b>95.62</b>	<b>98.49</b>	<b>96.48</b>
QB	6.11	15.53	<b>39.59</b>	22.46	33.63	<b>59.24</b>	42.57	49.83	69.81	43.64	48.32	68.88
T5	10.03	24.57	33.55	32.77	48.06	52.11	48.10	57.19	60.75	45.39	52.61	57.59

Table 18: Results of various approaches on the Chat-Ghosting task, on the DD dataset for different prefix lengths. Metrics at max TR and reported for seen test set.

	Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars		
	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec
MPC	<b>5.53</b>	<b>9.05</b>	25.61	<b>38.42</b>	<b>42.76</b>	<b>46.93</b>	<b>84.49</b>	<b>86.01</b>	<b>86.55</b>	<b>99.26</b>	<b>99.45</b>	<b>99.43</b>
MPC++	<b>5.53</b>	<b>9.05</b>	25.61	<b>38.42</b>	<b>42.76</b>	<b>46.93</b>	<b>84.49</b>	<b>86.01</b>	<b>86.55</b>	<b>99.26</b>	<b>99.45</b>	<b>99.43</b>
QB	2.57	6.51	<b>32.50</b>	5.34	11.79	44.85	7.20	14.08	52.94	7.35	13.56	51.32
T5	1.89	6.35	22.75	3.14	9.56	23.87	4.71	11.66	26.17	5.72	13.23	26.86

Table 19: Results of various approaches on the Chat-Ghosting task, on the DU dataset for different prefix lengths. Metrics at max TR and reported for seen test set.

	Contextual DD												Contextual DU														
	Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars			Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars					
	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec
MPC	0.01	4.64	23.01	0.06	7.47	18.24	0.66	11.89	16.64	8.17	<b>29.71</b>	29.88	0.04	3.14	20.09	0.19	6.21	13.03	0.33	8.50	11.19	0.00	<b>14.58</b>	17.18			
MPC++	0.01	4.64	23.01	0.39	7.61	19.24	4.09	12.18	27.35	9.08	15.92	34.84	0.04	3.15	20.09	0.25	6.17	14.47	1.59	8.10	23.41	2.45	9.17	31.12			
QB	0.14	3.48	<b>36.53</b>	1.16	5.55	<b>42.48</b>	4.70	9.87	<b>44.37</b>	8.15	13.50	<b>44.57</b>	<b>0.46</b>	<b>3.66</b>	<b>32.36</b>	1.50	6.74	<b>40.94</b>	3.47	8.94	<b>47.69</b>	<b>5.11</b>	10.13	<b>46.78</b>			
T5	<b>0.51</b>	<b>7.75</b>	20.53	<b>2.30</b>	<b>11.11</b>	21.87	<b>6.66</b>	<b>16.10</b>	25.03	<b>9.59</b>	18.90	27.53	0.41	<b>5.79</b>	21.46	<b>1.62</b>	<b>9.05</b>	22.16	<b>3.74</b>	<b>11.84</b>	25.37	4.72	12.91	26.06			

Table 20: Results of various approaches on the Chat-Ghosting task, on the contextual DD and contextual DU dataset for different prefix lengths. Metrics at max TR and reported for unseen test set.

	Contextual DD												Contextual DU														
	Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars			Len=1-5 chars			Len=6-12 chars			Len=13-25 chars			Len=26-50 chars					
	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec	MR	P-Rec	P-Prec
MPC	13.33	25.93	35.87	39.23	54.78	58.18	<b>70.90</b>	<b>83.18</b>	<b>77.52</b>	<b>96.53</b>	<b>99.05</b>	<b>97.25</b>	5.78	9.28	25.90	<b>40.07</b>	<b>44.29</b>	<b>48.62</b>	<b>85.16</b>	<b>86.66</b>	<b>87.26</b>	<b>99.25</b>	<b>99.40</b>	<b>99.42</b>			
MPC++	13.33	25.93	35.87	39.23	54.78	58.18	<b>70.90</b>	<b>83.18</b>	<b>77.52</b>	<b>96.53</b>	<b>99.05</b>	<b>97.25</b>	5.78	9.28	25.90	<b>40.07</b>	<b>44.29</b>	<b>48.62</b>	<b>85.16</b>	<b>86.66</b>	<b>87.26</b>	<b>99.25</b>	<b>99.40</b>	<b>99.42</b>			
QB	6.25	15.54	39.66	21.95	33.23	<b>58.78</b>	41.40	49.17	68.68	42.00	47.53	67.63	3.24	6.91	<b>35.95</b>	6.05	12.28	47.65	8.16	14.67	56.43	8.33	14.15	54.85			
T5	<b>19.60</b>	<b>36.35</b>	<b>40.52</b>	<b>36.97</b>	<b>55.53</b>	55.99	50.10	60.78	61.93	48.22	56.28	59.79	4.24	<b>10.23</b>	24.99	6.17	13.59	25.89	8.69	16.59	29.70	8.89	17.06	29.85			

Table 21: Results of various approaches on the Chat-Ghosting task, on the contextual DD and contextual DU dataset for different prefix lengths. Metrics at max TR and reported for seen test set.

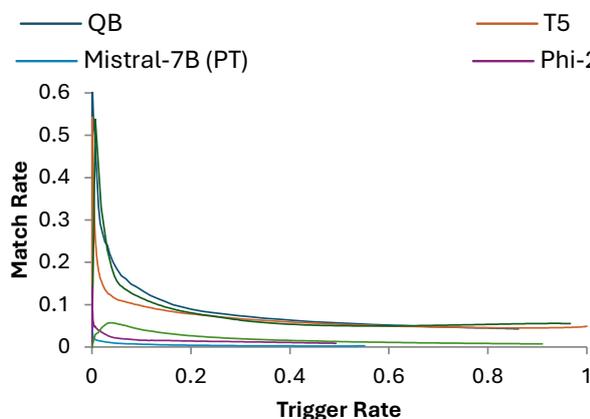


Figure 6: MR for DU

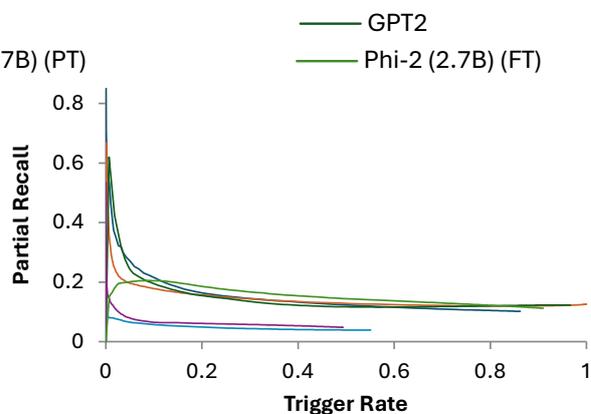


Figure 7: *P-Rec* for DU

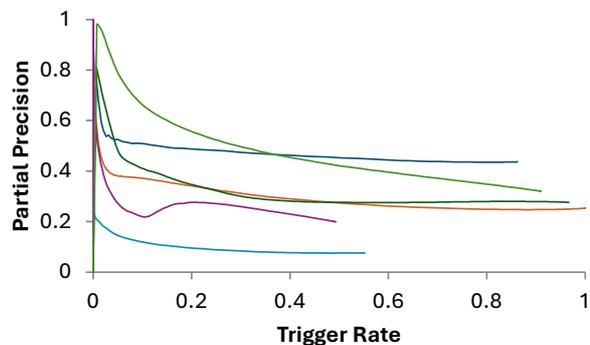


Figure 8: *P-Prec* for DU

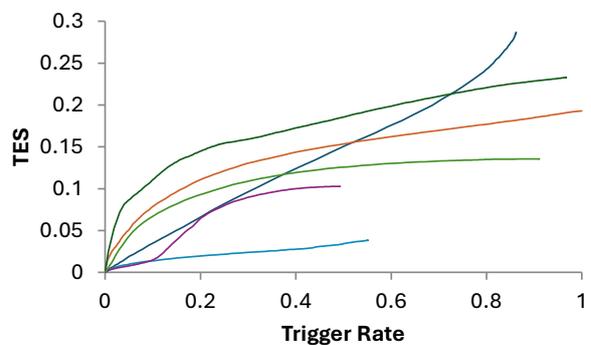


Figure 9: TES for DU

and TES should be higher for shorter suggestions. To verify this we truncate the predicted suggestions to  $t$  words where we vary  $t$  from 1 to 10. We show the results in Figs. 10 to 17 for the unseen test sets for the “without context” setting for both the DD and the DU datasets.

For match rate, typically there is high match rate with truncation set to 1 word and then it suddenly drops when truncation is done with 2 words. Post that the match rate continues to increase with trun-

cation length. We believe that the match rate is high for the 1 word case because it is much easier to complete the word that the user is currently typing compared to predicting full next few words.

As expected, *P-Prec* and TES reduce as we increase the truncation length. Also *P-Rec* increases as we increase the truncation length. Overall, these trends hold across both the datasets.

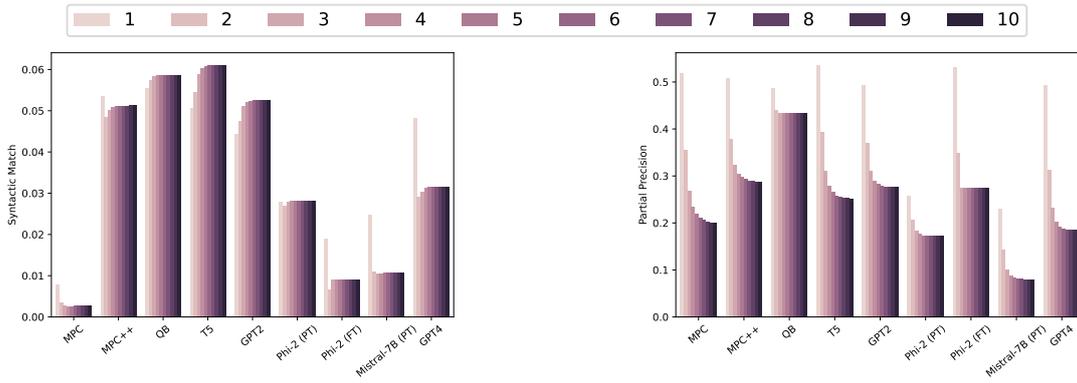


Figure 10: Truncation: MR for DD

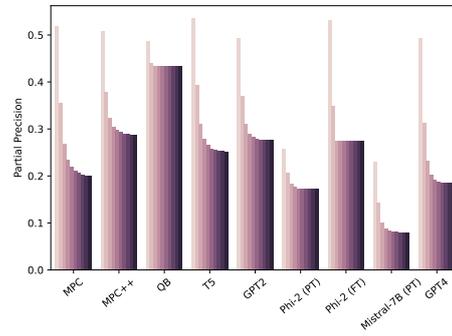


Figure 11: Truncation: *P-Prec* for DD

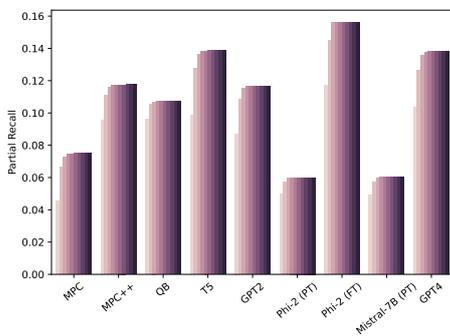


Figure 12: Truncation: *P-Rec* for DD

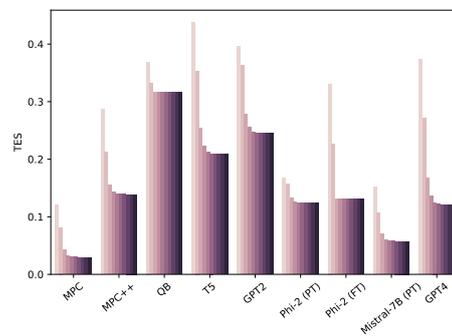


Figure 13: Truncation: TES for DD

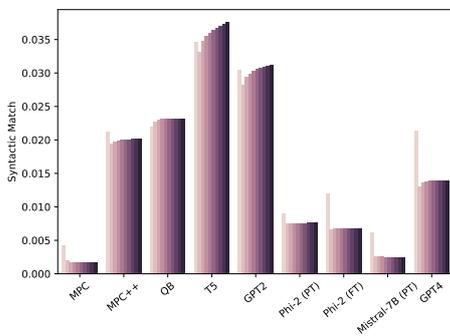


Figure 14: Truncation: MR for DU

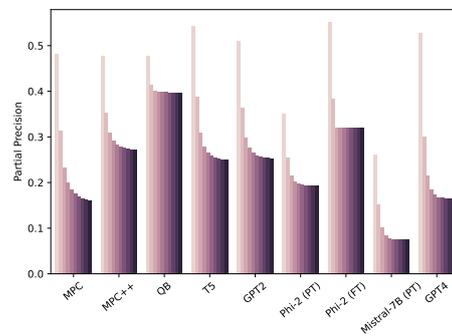


Figure 15: Truncation: *P-Prec* for DU

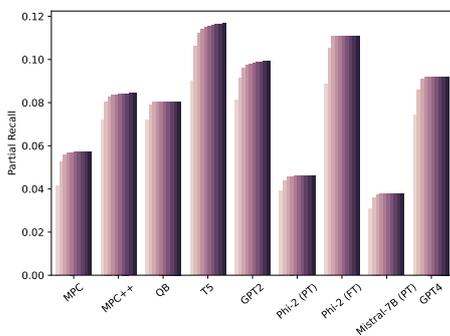


Figure 16: Truncation: *P-Rec* for DU

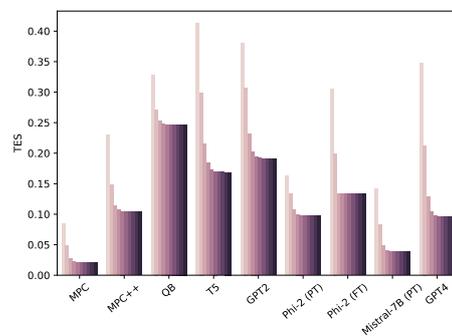


Figure 17: Truncation: TES for DU

Method	DD	DU	cDD	cDU	Hardware
<b>QB</b>	3	5	80	109	Intel Xeon Gold 6126 CPU
<b>MPC</b>	34	55	99	132	
<b>MPC++</b>	32	62	105	155	
<b>T5 (220M)</b>	85	86	103	132	1 V100 (32GB) with batch size=1
<b>GPT2 (124M)</b>	63	64	77	90	
<b>Phi-2 (2.7B) (PT)</b>	1605	1679	1711	1782	
<b>Phi-2 (2.7B) (FT)</b>	367	372	478	485	
<b>Mistral-7B (PT)</b>	2043	2381	2315	2647	

Table 22: Inference Latency for each chat-ghosting method (in msec) at max TR. PT=Pretrained, FT=Finetuned. cDD and cDU are Contextual datasets.

## L Optimizing Latency in Neural Language Model-Based Ghosting

Table 22 shows inference latency for various methods (in ms) at max TR. Note that Phi-2 finetuned latency is much smaller than the pretrained model. This is because for the pretrained models, we use 4-shot in-context learning, and generation token length is set to 40 tokens. For the Phi-2 finetuned model, we finetune using just the  $\langle$ context, prefix, suffix $\rangle$  without any exemplars resulting into a much shorter input. Also, max generation length for Phi-2 finetuned model is set to 5. As expected, adding context increases latency.

We can observe that inference can be done using models like T5 and GPT2 within 100-150ms on a V100 GPU with batch size=1. While this may be within latency budget for practically deployed web scale models, further optimizations can be done to reduce latency and increase throughput. Throughput is also important when catering to millions of users with limited GPU investment. Latency with Phi-2 and Mistral models is much larger (1-3 seconds) and clearly needs to be optimized further for practical deployments.

While our current focus was on evaluating the fundamental performance characteristics of various modeling strategies including utilization of SLMs or LLMs, we recognize the crucial importance of techniques like model pruning, distillation, and caching for practical deployment in real-time scenarios (Gupta and Agrawal, 2022). We could perform quantization aware training or post-training quantization using methods like GPTQ (Frantar et al., 2022) and AWQ (Activation Aware Quantization) (Lin et al., 2024) to quantize models to 4, 8 or 16 bits. We could also perform CUDA-level optimizations like ONNX optimization (Jin et al., 2020) (which builds a universal CUDA graph and fuses different submodule operations), Tensor-RT LLM (NVIDIA, 2025) (which supports quantization to 4/8 bit

weights and 16 bit activations for Nvidia GPUs), vLLM (with Paged-attention) (Kwon et al., 2023), FasterTransformer (Chelba et al., 2020), Deepspeed FastGen (Holmes et al., 2024), and Llama.cpp (Gerganov, 2025). Of course, pruning the models is an orthogonal direction to explore and previous research (Han et al., 2015) has shown that pruning and quantization are complementary. Lastly, we could distill the large neural models to their smaller pretrained variants using popular knowledge distillation methods (Gupta and Agrawal, 2022).

For faster decoding, there are a large number of heuristics that could be leveraged like FasterTransformer (Chelba et al., 2020), setting a maximum prefix length, reducing beam width, and performing FP16 decoding. We could also batch requests to increase throughput. But batching could lead to slight increase in latency. A dynamic batch size could strike a good tradeoff by providing high throughput during high traffic periods and low latency during low traffic periods.

## M Semantic Evaluation

Our study mainly focuses on syntactic metrics to assess the quality of various ghosting methods. However, several predictions may not have syntactic match but they could be semantically meaningful. Hence, we conducted human evaluation of a set of 100 instances each from seen and unseen test set of the DD dataset. The evaluation was done for outputs from each of our 10 methods. We also obtained GPT4o evaluation for the same set. We note that GPT4o evaluation has a correlation of 0.57 with human evaluation with p-value of  $2.2e-88$ , suggesting that GPT4o can be used as a replacement for human evaluation on the entire set.

We also calculated METEOR score as a measure for semantic similarity. METEOR’s effectiveness as a metric was demonstrated by the strong correlation between its ranking of suggestions and the ranking produced by GPT-4o, as measured by the length of the Longest Common Subsequence (LCS).

We used the following GPT-4o prompt for evaluation.

```
<lim_start>system
[system](#instructions)
# ChatGhosting: Evaluating Auto-Completion Systems for
Chatbots
```

The primary purpose of this human evaluation is to assess

the quality and effectiveness of chat autocompletion models.

### ## Instructions

Your task is to evaluate how helpful and relevant the autocompletion predictions are by rating them on a integer scale of 1 to 5.

Imagine you are actually typing the message and considering these suggestions.

- 1: Bad completion, so not helpful or relevant.
- 2: Some other completion, not helpful or relevant.
- 3: Not exactly the same intent, but still helpful or relevant.
- 4: Doesn't match the ground truth but means the same thing, so user can still use it.
- 5: Matches the ground truth exactly.

Please rate the following completions based on the above criteria.

### ## Format

For each candidate response, provide a rating from 1 to 5. Use json format to provide the ratings.

Along with each rating, provide a brief explanation of why you chose that rating.

### ## Examples

```
```json
{
  "Model 1": {
    "rating": 4,
    "reason": "The completion is not exactly the same but has the same intent."
  }
  "Model 2": {
    "rating": 5,
    "reason": "The completion matches the ground truth exactly."
  }
  ...
  "Model 10": {
    "rating": 3,
    "reason": "The completion is not exactly the same intent but still helpful."
  }
}
```
```

<!--end-->

<!--start-->user

## Input Context

Context: #context#

Ground Truth: #ground\_truth#

## User Input

Prefix Typed by User: #prefix#

## Candidate AutoCompletions

Model 1: #model\_1#

Model 3: #model\_3#

Model 5: #model\_5#

Model 7: #model\_7#

Model 9: #model\_9#

Model 11: #model\_11#

Model 2: #model\_2#

Model 4: #model\_4#

Model 6: #model\_6#

Model 8: #model\_8#

| Model        | Human Scores | Eval | GPT4 Scores | Eval | Corr (Human, GPT4) | pval (Human, GPT4) |
|--------------|--------------|------|-------------|------|--------------------|--------------------|
| Context+GPT2 | 2.85         |      | 2.06        |      | 0.57               | 4.29E-10           |
| T5           | 2.28         |      | 1.92        |      | 0.63               | 1.61E-12           |
| GPT2         | 2.255        |      | 1.76        |      | 0.60               | 5.82E-11           |
| QB           | 2.215        |      | 1.81        |      | 0.36               | 3.07E-05           |
| Context+QB   | 2.165        |      | 1.81        |      | 0.40               | 3.07E-05           |
| Context+T5   | 2.115        |      | 1.68        |      | 0.56               | 9.47E-10           |
| MPC          | 2.045        |      | 1.54        |      | 0.66               | 4.73E-14           |
| MPC++        | 2.04         |      | 1.54        |      | 0.66               | 4.67E-14           |
| Context+MPC  | 2.025        |      | 1.54        |      | 0.68               | 9.93E-15           |

Table 23: Detailed human evaluation results on the 100 instances

### Purpose

The primary purpose of this human evaluation is to assess the quality and effectiveness of chat autocompletion models.

### Objective

Your task is to evaluate how helpful and relevant the autocompletion predictions are by rating them on an integer scale of 1 to 5. Imagine you are typing the message and considering these suggestions.

### Definitions

**Context:** Shows the previous chat utterances. Different speakers' utterances are separated by the <!-- token.

**Prefix:** Part of the current utterance you've already typed.

**Ground\_truth\_suffix:** The actual continuation of the prefix that the original user typed.

**Predictions:** There are 5 different models. Each model gives 2 predictions: one using only the "Prefix" (non-contextual) and one using both "Prefix" and "Context" (contextual).

### Evaluation

Read the Context, Prefix, and Ground Truth Suffix: Understand the situation and the intended message.

Examine Each Prediction: For each of the 10 predictions (5 models x 2 context variations), consider:

- 1: Bad completion, so not helpful or relevant.
- 2: Some other completions, not helpful or relevant.
- 3: Not exactly the same intent, but still helpful or relevant.
- 4: Doesn't match the ground truth but means the same thing, so the user can still use it.
- 5: Match the ground truth exactly.

Figure 18: Human Evaluation Task Snapshot

Model 10: #model\_10#  
Model 12: #model\_12#

```
<!--end-->
<!--start-->assistant
[assistant]#evaluation)
```json
```

Correspondingly, for human evaluation the task given to the annotators is shown in Fig. 18. Table 23 shows the detailed human evaluation results on the 100 instances. This table shows moderate and statistically significant correlations with GPT-4o-based evaluations for all models ( $r = 0.56-0.68$  for the top methods, all  $p < 1e-9$ ). We observe that context helps overall from a semantic match perspective. Thus, human evaluation, GPT4o evaluation, and semantic metrics suggest that contextual models are better than non-contextual models. This is also a trend we observe from automated metrics. This confirms that our automatic metrics are reliable and that the reported trends are not artifacts of metric choice.

Hence, we evaluate all models on the entire test

Rank	Unseen (GPT4)	Seen (GPT4)	Seen (METEOR)	Unseen (METEOR)
1	Context+T5	Context+MPC	Context+MPC++	Context+T5
2	Context+GPT2	Context+MPC++	Context+MPC	T5
3	T5	MPC	MPC	Context+GPT2
4	QB	MPC++	MPC++	GPT2
5	Context+QB	Context+T5	Context+T5	QB
6	GPT2	T5	T5	Context+MPC++
7	Context+MPC++	Context+GPT2	Context+GPT2	Context+QB
8	MPC++	GPT2	GPT2	MPC++
9	Context+MPC	QB	QB	Context+MPC
10	MPC	Context+QB	Context+QB	MPC

Table 24: Semantic evaluation of our Chat-Ghosting methods on DD test sets.

set of the DD dataset using both GPT4o-prompt and METEOR. We show the results in Table 24. The table shows the rankings of various methods. We observe that context helps overall from a semantic match perspective. Thus, human evaluation, GPT4o evaluation, and semantic metrics suggest that contextual models are better than non-contextual models.

## N Frequently asked questions

### N.1 Why did we mainly focus on prefix match based metrics?

Standard metrics (BLEU, ROUGE, METEOR, perplexity) are not directly applicable for our task because: (1) it is important for us to match prefix (first few chars) of the suffix. Std metrics give equal importance to all words. (2) Word order is important for our task. Standard metrics ignore word order. (3) Generation of exact ground truth (and not just assigning high probability to those tokens) is important.

Semantic match based metrics like METEOR and BERTScore are becoming popular for various natural language generation tasks. However, this does not hold for Autocompletion systems. In Autocompletion, the user typically has a completion in mind, and they are splitting their cognitive bandwidth between typing and observing matching suggestions. Doing a semantic match with the shown suggestions leads to high cognitive load. Hence, if the suggestions are not syntactic exact matches, highly likely the user will not click. Hence, we mainly use syntactic prefix matching based metrics, rather than semantic matching ones.

However, for sake of completeness we also reported results using METEOR and GPT-4o evaluation in Appendix M.

### N.2 Why did we use DialogCC, DSTC7-Ubuntu, Open Assistant and ShareGPT datasets?

We have evaluated our proposed methods using one task-specific dataset (DSTC7-Ubuntu) and one general open-ended conversations dataset (DialogCC). We chose DSTC7-Ubuntu because it is a standard dataset widely used in the dialog modeling literature. Also, DSTC7-Ubuntu is focused on Ubuntu technical domain, and we chose it to assess the degree to which our methods can handle technical difficulties of autocompleting code compared to other tasks. Further, DialogCC is a general open-ended conversations dataset and its purpose is to judge the performance of the model (and its generalizability) across dialogues of different domains and genres.

Although DD and DU are popular benchmark datasets in the dialog modeling community, they involve human-human conversations only. We also wanted to test our methods with human-bot conversations. Hence, we also experiment with two other datasets: Open Assistant and ShareGPT. The bot utterances are typically much larger. We have taken care to generate train and test samples from the human utterances only. Open Assistant is a human-annotated assistant conversation corpus. ShareGPT contains user-LLM-chatbots conversations collected by the ShareGPT API.

### N.3 Why is TES a useful metric?

TES literally captures number of keystrokes saved which is the main goal of autocompletion systems. While other metrics are at prefix level, TES is computed at utterance level. Interestingly, sometimes TES may decrease even when MR increases. Hence, a good system is one which has both high MR as well as high TES. E.g., consider the query abcde where system A generates x (some mismatching completion) as completion for prefixes a, abc, abcd and abcde. For prefix=ab, it generates cde as suffix. Here, MR=1/4, TES=3/5. Consider system B that generates x (some mismatching completion) as completion for prefixes a and ab. For prefix=abc and abcd, it generates de and e as suffix respectively. Here, MR=2/4 but TES=2/5. System A has higher TES but lower MR compared to system B.

#### **N.4 Why not have a metric to measure impact due to incorrect suggestions?**

Incorrect suggestions can indeed degrade the user experience, and any robust metric should ideally account for this. The TES metric implicitly penalizes incorrect suggestions. This is because, by design, TES only considers characters saved when the suggested prefix is an actual prefix of the ground truth. If a suggestion is incorrect, the user will be forced to type additional character(s) to correct or complete their input. In these cases, there would be no reduction in keystrokes, which means that incorrect suggestions result in a lower TES value, thereby providing a form of penalty.

But incorrect suggestions could also lead to cognitive verification and suggestion-rejection load. TES does not explicitly capture the cognitive burden associated with examining suggestions. However, objectively quantifying this cognitive load is a significant challenge without access to user interaction data. The user typing characteristics, like slow typing, slow reading and examining will highly influence the cognitive load experienced. Therefore, we chose to focus on the directly measurable aspect of typing effort reduction as a more tractable starting point for our research.

In summary, while TES provides a valuable measure of keystroke reduction and inherently penalizes incorrect suggestions, we recognize its limitations in capturing cognitive load. Our future work aims to explore methods for refining metrics to better address this challenge.

#### **N.5 Is it possible to use token based trie methods instead of character based trie methods for better data compression and potentially faster inference time?**

Token-based tries are theoretically possible, but they do not offer practical gains for chat-ghosting. Token boundaries misalign with keystrokes, causing frequent partial-token states that must be mapped back to character prefixes, which removes most of the compression benefit. In our experiments, the overhead of handling partial tokens outweighed any speedup from token-level storage. Character-based tries therefore remain simpler and more reliable for per-keystroke generation.