

(Chat)GPT v BERT Dawn of Justice for *Semantic Change Detection*

Francesco Periti	Haim Dubossarsky	Nina Tahmasebi
University of Milan	Queen Mary University of London	University of Gothenburg
Via Celoria 18	Mile End Road	Renströmsgatan 6
20133 Milan, Italy	E1 4NS London, United Kingdom	40530 Göteborg, Sweden
francesco.periti@unimi.it	h.dubossarsky@qmul.ac.uk	nina.tahmasebi@gu.se

Accepted to the Findings of EACL 2024

Abstract

In the *universe* of Natural Language Processing, Transformer-based language models like BERT and (Chat)GPT have emerged as *lexical superheroes* with *great power* to solve open research problems. In this paper, we specifically focus on the temporal problem of semantic change, and evaluate their ability to solve two diachronic extensions of the Word-in-Context (WiC) task: TempoWiC and HistoWiC. In particular, we investigate the potential of a novel, off-the-shelf technology like ChatGPT (and GPT) 3.5 compared to BERT, which represents a family of models that currently stand as the state-of-the-art for modeling semantic change. Our experiments represent the first attempt to assess the use of (Chat)GPT for studying semantic change. Our results indicate that ChatGPT performs significantly worse than the foundational GPT version. Furthermore, our results demonstrate that (Chat)GPT achieves slightly lower performance than BERT in detecting long-term changes but performs significantly worse in detecting short-term changes.

1 Introduction

Lexical semantic change is the linguistic phenomenon that denotes words changing their meanings over time (Geeraerts et al., 2024; Bloomfield, 1933). An example is the word *gay* that changed from meaning *cheerful* to *homosexual* in the last century. This change is crucial to our understanding of historical texts. A nuanced grasp of semantic *variation* between groups and genre, and semantic *change* across time allows us to study languages, cultures, and societies through digitized text and opens up a range of research applications. Computational approaches to semantic change are thus tools with immense potential for a range of research fields (Montanelli and Periti, 2023; Tahmasebi et al., 2021; Kutuzov et al., 2018). Not only can they broaden the field of historical linguistics and simplify lexicography, but they can

also be fruitfully applied in the fields of sociology, history, and other text-based research. For instance, the computational modeling of semantic change is equally relevant when studying out-of-domain texts where language differs from the general language, like in medical (Kay, 1979) and olfactory (Paccosi et al., 2023; Menini et al., 2022) domains.

The recent introduction of Transformer-based (Vaswani et al., 2017) language models (LMs) has led to significant advances in Natural Language Processing (NLP). These advances are exemplified in Pretrained Foundation Models like BERT (Devlin et al., 2019) and GPT, which “are regarded as the foundation for various downstream tasks” (Zhou et al., 2023). BERT has experienced a surge in popularity over the last few years, and the family of BERT models has repeatedly provided state-of-the-art (SOTA) results for computational modeling of semantic change (Cassotti et al., 2023; Periti et al., 2023). However, research focus is now shifting toward ChatGPT due to its impressive ability to generate fluent and high-quality responses to human queries, making it the fastest-growing AI tool. Several recent research studies have assessed the language capabilities of ChatGPT by using a wide range of prompts to solve popular NLP tasks (Laskar et al., 2023; Kocoń et al., 2023). However, current evaluations generally (a) overlook the fact that the output of ChatGPT is nondeterministic,¹ (b) rely only on contemporary and synchronic text, and (c) consider predictions generated by the ChatGPT² web interface, which is based on the Chat version of the GPT foundation model. As a result, these evaluations provide valuable insights into the generative, pragmatic, and semantic capabilities of ChatGPT (Kocoń et al., 2023), but fall short when it comes to assess the potential of GPT to solve

¹platform.openai.com/docs/guides/gpt/faq

²chat.openai.com



Figure 1: The title of this paper draws inspiration by the movie *Batman v Superman: Dawn of Justice*. We leverage the analogy of (Chat)GPT and BERT, powerful and popular LMs, as two lexical superheroes often erroneously associated for solving similar problems. Our aim is to shed lights on the potential of (Chat)GPT for semantic change detection.

NLP tasks and specifically to handle historical and diachronic text, which constitutes a unique scenario for testing models’ ability to generalize.

In this paper, we propose to evaluate the use of both ChatGPT and GPT - i.e., (Chat)GPT³ - to recognize (lexical) semantic change. Our goal is not to comprehensively evaluate (Chat)GPT in dealing with semantic change but rather to evaluate its potential as *off-the-shelf* model with a *reasonable* prompts from a human point of view, which may not necessarily be optimized for the model. Recently, a novel evaluation task in NLP, called Lexical Semantic Change (LSC), has been introduced as a shared task at SemEval (Schlechtweg et al., 2020). The LSC task involves considering all occurrences (potentially several thousands) of a set of target words to assess their change in meaning within a diachronic corpus. As a result, this setup is *currently* not suitable for evaluating a GPT model, due to the limited size of its prompts and answers, as well as accessibility limitations such as an hourly character limit and economic constraints. In light of these considerations, we chose to evaluate the potential of (Chat)GPT through the Word-in-Context (WiC, Pilehvar and Camacho-Collados, 2019) task, which has recently demonstrated a robust connection with LSC (Cassotti et al., 2023; Arefyev et al., 2021). In particular, we consider two diachronic extensions of the original WiC setting, namely *temporal* WiC (TempoWiC, Loureiro et al., 2022) and *historical* WiC (HistoWiC). Our

goal is to determine whether a word carries the same meaning in two different contexts of different time periods, or conversely, whether those contexts exemplify a semantic change. While TempoWiC has been designed to evaluate LMs ability to detect short-term changes in social media, HistoWiC is our adaptation of the SemEval benchmark of historical text to a WiC task for evaluating LMs ability to detect long-term changes in historical corpora.

Considering the remarkable performance of contextualized BERT models in addressing WiC and LSC tasks (Montanelli and Periti, 2023; Periti and Dubossarsky, 2023; Periti et al., 2023), we compare the performance of (Chat)GPT in TempoWiC and HistoWiC to those obtained using BERT. While BERT is specifically designed to understand the meaning of words in context, (Chat)GPT is designed to generate fluent and coherent text. Through these two *lexical superheroes* (see Figure 1), we aim to illuminate the potential of (Chat)GPT as *off-the-shelf* model and mark *the dawn of a new era* by assessing whether it *already* makes the approaches to WiC and LSC, which rely on BERT-embedding similarities, outdated.

2 Related work

The significant attention garnered by ChatGPT has led to a large number of studies being published immediately after its release. Early studies mainly focused on exploring the benefits and risks associated with using ChatGPT in expert fields such as education (Lund and Wang, 2023), medicine (Antaki et al., 2023), or business (George and George, 2023). Evaluation studies are currently emerging for assessing (Chat)GPT’s generative and linguistic capabilities across a wide range of downstream tasks in both monolingual and multilingual setups (Bang et al., 2023; Shen et al., 2023; Lai et al., 2023). Most evaluations focus on ChatGPT and involve a limited number of instances (e.g., 50) for each task considered (Weissweiler et al., 2023; Zhong et al., 2023; Alberts et al., 2023; Khalil and Er, 2023). When the official API is used to query the GPT foundation model, this limit is imposed by the hourly token processing limit⁴ and the associated costs.⁵ When the web interface is used instead of the API, the limit is due to the time-consuming

³Throughout the text, we distinguish between ChatGPT, which is the standard (web) version of GPT, and GPT, which serves as the foundation model. Instances of (Chat)GPT represent both types of models.

⁴help.openai.com/en/articles/4936856-what-are-tokens-and-how-to-count-them

⁵openai.com/pricing

process of interacting with ChatGPT that keeps humans in the loop. Thus far, even systematic and comprehensive evaluations (Kocoń et al., 2023; Laskar et al., 2023) rely on repetition of a single experiment for each task. However, while individual experiments provide valuable insights into (Chat)GPT’s capabilities, they fall short in assessing the potential of (Chat)GPT to solve specific tasks given its nondeterministic nature. Multiple experiments need to be conducted to validate its performance on each task. In addition, current evaluations generally leverage tasks that overlook the temporal dimension of text, leaving a gap in our understanding of (Chat)GPT’s ability to handle diachronic and historical text.

Our original contribution.

Our evaluation of (Chat)GPT focuses on two diachronic extensions of the WiC task, namely TempoWiC and HistoWiC. Our aim is to assess the potential of (Chat)GPT for **Semantic Change Detection**. To the best of our knowledge, this paper is the first to investigate the application of (Chat)GPT for historical linguistic purposes. Thus far, only the use of ChatGPT for a conventional WiC task has been evaluated by Laskar et al. (2023) and Kocoń et al. (2023), who reported low accuracy under a single setup. In this paper, we challenge their performance by considering diachronic text and the following setups, totaling 47 experiments each for TempoWiC and HistoWiC:

- **Different prompts.** Like Zhong et al. (2023), we evaluate (Chat)GPT using zero-shot and few-shot prompting strategies, while also exploring many-shot prompting. Our results demonstrate that zero-shot prompting is more effective on HistoWiC, while few-shot prompting is more effective on TempoWiC.
- **Varying temperature.** Like Peng et al. (2023); Liu et al. (2023), we analyze how GPT’s performance varies according to its temperature hyperparameter, which controls the “creativity” or randomness of its answers. Our results indicate that GPT used with low temperature values (i.e., less creativity) is better at handling WiC tasks.
- **GPT API v ChatGPT Web.** We empirically assess whether GPT produces worse results through the OpenAI API compared to Chat-

GPT through the web interface.⁶ Our results demonstrate that using GPT through the official API for WiC tasks is better than using ChatGPT through the web interface, as has previously been done (Laskar et al., 2023; Kocoń et al., 2023). Furthermore, our findings suggest that the web interface automatically sets an intermediate temperature for ChatGPT.

- **(Chat)GPT v BERT.** Finally, like Zhong et al. (2023), we compare the performances of (Chat)GPT and BERT. By leveraging the TempoWiC task and introducing the novel HistoWiC task, we shed light on the potential of both models and demonstrate the *current* superiority of BERT in dealing with diachronic text and WiC tasks, compared to *reasonable* GPT prompts templates and strategies.

3 Semantic Change Detection

Our evaluation relies on two diachronic definitions of the conventional Word-in-Context (WiC) task, namely TempoWiC and HistoWiC. WiC is framed as a binary classification problem, where each instance is associated with a target word w , either a verb or a noun, for which two contexts, c_1 and c_2 , are provided. The task is to identify whether the occurrences of w in c_1 and c_2 correspond to the same meaning or not. Both TempoWiC and HistoWiC rely on the same definition of the task, while being specifically designed for semantic change detection in diachronic text.

3.1 Temporal Word-in-Context

NLP models struggle to cope with new content and trends. TempoWiC is designed as an evaluation benchmark to detect short-term semantic changes on social media, where the language is extremely dynamic. It uses tweets from different time periods as contexts c_1 and c_2 .

Given the limits on testing (Chat)GPT, we followed Zhong et al. (2023); Jiao et al. (2023) and randomly sampled a subset of the original TempoWiC datasets. While the original TempoWiC framework provides trial, train, test, and dev sets, here we did not consider the dev set. Table 1 shows the number of positive (i.e., same meaning) and

⁶Discussions on this topic are currently very active, for example, community.openai.com/t/web-app-vs-api-results-web-app-is-great-api-is-pretty-awful/96238

negative (i.e., different meanings due to semantic change) examples we considered for each set.

Table 1: Datasets used in our evaluation

	TempoWiC			HistoWiC		
	Trial	Train	Test	Trial	Train	Test
<i>True</i>	8	86	73	11	137	79
<i>False</i>	12	114	127	9	103	61
Total	20	200	200	20	200	140

3.2 Historical Word-in-Context

Given that NLP models also struggle to cope with historical content and trends, we designed HistoWiC as a novel evaluation benchmark for detecting long-term semantic change in historical text, where language may vary across different epochs. HistoWiC sets the two contexts, c_1 and c_2 , as sentences collected from the two English corpora of the LSC detection task (Schlechtweg et al., 2020).

Similar to the original WiC (Pilehvar and Camacho-Collados, 2019), the annotation process for the LSC benchmark involved usage pair annotations where a target word is used in two different contexts. Thus, we directly used the annotated instances of LSC to develop HistoWiC. Since LSC instances were annotated using the DUREL framework (Schlechtweg et al., 2023) and a four-point semantic-relatedness scale (Schlechtweg et al., 2021, 2020, 2018), we only binarized the human annotations (see Appendix A).

As with TempoWiC, we randomly sampled a limited number of instances to create trial, training, and test sets. Table 1 shows the number of positive and negative examples for each set.

4 Experimental setup

In the following, we present our research questions (RQs) and the various setups we considered in our work. In our experiments, we evaluated the performance of (Chat)GPT 3.5 over the TempoWiC and HistoWiC test sets using both the official OpenAI API (GTP API)⁷ and the web interface (ChatGPT Web).⁸ Of the GPT 3.5 models available through the API, we assessed the performance of gpt-3.5-turbo. Following Loureiro et al. (2022), we employed the Macro-F1 for multiclass classification problems as evaluation metric.

⁷version 0.27.8.

⁸The August 3 Version.

4.1 (Chat)GPT prompts

Current ChatGPT evaluations are typically performed manually (Laskar et al., 2023). When automatic evaluations are performed, they are typically followed by a manual post-processing procedure (Kocoń et al., 2023). As manual evaluation and processing may be biased due to answer interpretation, we addressed the following research question:

RQ1: *Can we evaluate (Chat)GPT in WiC tasks in a completely automatic way?*

Furthermore, as current evaluations generally rely on a zero-shot prompting strategy, we addressed the following research question:

RQ2: *Can we enhance (Chat)GPT’s performance in WiC tasks by leveraging its in-context learning capabilities?*

To address RQ1 and RQ2, we designed a prompt template to explicitly instruct (Chat)GPT to answer in accordance with the WiC label format (i.e., *True*, *False*). We then used this template (see Appendix C.1) with different prompt strategies:

- *zero-shot prompting* (ZSp): (Chat)GPT was asked to address the WiC tasks (i.e., test sets) without any specific training, generating coherent responses based solely on its preexisting knowledge.
- *few-shot prompting* (FSp): since PFMs have recently demonstrated *in-context learning* capabilities without requiring any fine-tuning on task-specific data (Brown et al., 2020), (Chat)GPT was presented with a limited number of input-output examples (i.e., trial sets) demonstrating how to perform the task. The goal was to leverage the provided examples to improve the model’s task-specific performance.
- *many-shot prompting* (MSp): similar to FSp, but with a greater number of input-output examples (i.e., training sets).

4.2 (Chat)GPT temperature

The temperature is a hyperparameter of (Chat)GPT that regulates the variability of responses to human queries. According to the OpenAI FAQ, the temperature parameter ranges from 0.0 to 2.0, with lower values making outputs mostly deterministic

and higher values making them more random.⁹ To counteract the nondeterminism of (Chat)GPT, we focused only on TempoWiC and HistoWiC and conducted the same experiment multiple times with progressively increasing temperatures. This approach enabled us to answer the following research questions:

RQ3: *Does (Chat)GPT demonstrate comparable effectiveness in detecting short-term changes in contemporary text and long-term changes in historical text?*

RQ4: *Can we enhance (Chat)GPT’s performance in WiC tasks by raising the “creativity” using the temperature value?*

To address RQ3 and RQ4, we evaluated GPT API in TempoWiC and HistoWiC using eleven temperatures in the range [0.0, 2.0] with 0.2 increments. For each temperature and prompting strategy, we performed two experiments and considered the average performance.

4.3 GPT API v ChatGPT Web

Current evaluations typically prompt GPT through the web interface instead of the official OpenAI API. This preference exists because the web interface is free and predates the official API. However, there are differences between using ChatGPT through the web interface (ChatGPT Web) and the official API (GPT API). First of all, the official API enables to query the GPT foundation model, while the web interface enables to query the Chat version. In addition, the GPT API can be set to test at varying temperatures, but the temperature value on ChatGPT Web cannot be controlled. However, while the GPT API allows a limited message history, ChatGPT Web seems to handle an unlimited message history (see Appendix B).

We used the following research question to compare the performance of GPT API and ChatGPT Web:

RQ5: *Does GPT API demonstrate comparable performance to ChatGPT Web in solving WiC tasks?*

Testing GPT API with the MSp strategy would be equivalent to testing it with the FSp strategy due to the limited message history. Thus, we evaluated

ChatGPT Web with MSp, aiming to address the following research question:

RQ6: *Can we enhance ChatGPT’s performance in WiC tasks by providing it with a larger number of in-context examples?*

To address these research questions, we tested (Chat)GPT using a single chat for each prompting strategy considered. Since testing ChatGPT Web is extremely time-consuming, we conducted one experiment for each prompting strategy.

4.4 (Chat)GPT v BERT

The ability of (Chat)GPT to understand has prompted the belief that ChatGPT is a *jack of all trades* that makes previous technologies somewhat outdated. Drawing upon Kocoń et al. (2023), we believe that, when used for solving downstream tasks as *off-the-shelf* model, (Chat)GPT is *currently* a *master of none*. It works on a comparable level to the competition, but does not outperform any major SOTA solutions.

By relying on multiple experiments on TempoWiC and HistoWiC, we aimed to empirically assess the potential of (Chat)GPT for WiC and LSC tasks. In particular, we addressed the following research question:

RQ7: *Does (Chat)GPT outperform BERT embeddings in detecting semantic changes?*

To address RQ7, we evaluated bert-base-uncased on TempoWiC and HistoWiC over different layers. Recent research has exhibited better results when utilizing earlier layers rather than the final layers for solving downstream tasks such as WiC (Periti and Dubossarsky, 2023; Ma et al., 2019; Reif et al., 2019; Liang and Shi, 2023). For each layer, we extracted the word embedding for a specific target word w in the context c_1 and c_2 . Since the focus of our evaluation was on (Chat)GPT, we did not fine-tune BERT and simply used the similarity between the embeddings of w in the context c_1 and c_2 . In particular, we followed Pilehvar and Camacho-Collados (2019), and trained a threshold-based classifier using the cosine distance between the two embeddings of each pair in the training set. The training process consisted of selecting the threshold that maximized the performance on the training set. We trained a distinct threshold-based classifier for each BERT layer and for each WiC task (i.e., TempoWiC

⁹platform.openai.com/docs/api-reference/chat

and HistoWiC). Then, in our evaluation, we applied these classifiers to evaluate BERT over the TempoWiC and HistoWiC test sets.

Finally, we addressed the following research question:

RQ8: *Can we rely on the pretrained knowledge of GPT to automatically solve the LSC task?*

Since (Chat)GPT has demonstrated awareness of historical lexical semantic changes when manually asked about the lexical semantic changes of some words (e.g., *plane*), our goal with RQ8 was to automatically test GPT’s pretrained knowledge of historical semantic changes covered in the English LSC benchmark. In addressing this research question we relied on the LSC ranking task as defined in [Schlechtweg et al. \(2018\)](#). Thus, we specifically asked GPT to rank the set of 37 target words in the English LSC benchmark according to their degree of LSC between two time periods, T1 (1810–1860) and T2 (1960–2010). For each temperature, we repeated the same experiment ten times, totaling 110 experiments. Then, for each temperature, we evaluated GPT’s performance by computing the Spearman correlation using gold scores derived from human annotation and the average GPT score for each target (see Appendix C.2).

5 Experimental results

In this section, we report the results of our experiments, while discussing the findings in regard to each research question.¹⁰

RQ1: (Chat)GPT consistently followed our template in nearly all cases, thereby allowing us to evaluate its answers without human intervention. For GPT API, however, we noticed that the higher the temperature, the larger the tendency for deviations from the expected response format (see Figure 2). ChatGPT Web only once answered with an incorrect format. To ensure impartiality, we classified the few (Chat)GPT responses that did not adhere to the required format as incorrect answers.

RQ2: Figure 3 shows the rolling average of the performance of GPT API across different temperatures, prompting strategies, and WiC tasks. By using a window size of 4, we were able to consider 8 different experiments per temperature (for each




Figure 2: Average number of wrongly formatted answers (WFAs) over the temperature values considered. Background lines correspond to each experiment.




Figure 3: Performance of GPT API (Macro-F1) as temperature increases.




Figure 4: Performance of ChatGPT Web (Macro-F1). Temperature is unknown.

temperature, we ran two experiments)¹¹. Figure 4 shows the performance of ChatGPT Web across different prompting strategies and WiC tasks.

Figure 3 and 4 show that ZSp consistently outperforms FSp on HistoWiC. By contrast, FSp consistently outperforms ZSp in TempoWiC when the GPT API is used. This result suggests that the in-context learning capability of GPT is more limited for historical data. In Figure 4, ChatGPT Web’s performance with ZSp outperforms that obtained with FSp for both TempoWiC and HistoWiC, although the discrepancy is smaller.

RQ3: Figures 3 and 4 show that (Chat)GPT’s performance on TempoWiC is consistently lower than its performance on HistoWiC. In particular, in our

¹⁰We provide all our data, code, and results at <https://github.com/FrancescoPeriti/ChatGPTvBERT>

¹¹Except for the first and last two temperatures.

Table 2: Macro-F1 scores obtained by SOTA systems, (Chat)GPT (best score), and BERT (last layer).

	Macro-F1
Chen et al., 2022	.770
Loureiro et al., 2022	.703
Loureiro et al., 2022	.670
Lyu et al., 2022	.625
<i>GPT API</i>	.689
<i>ChatGPT Web</i>	.580
<i>BERT</i>	.743

experiments we observe that (Chat)GPT’s performance ranges from .551 to .689 on TempoWiC and from .552 to .765 on HistoWiC. This suggests that (Chat)GPT is significantly more effective for long-term change detection than for short-term change detection. We believe that this might involve word meanings that were not explicitly covered during training, potentially allowing (Chat)GPT to detect anomalies from the usual patterns. We will further investigate this aspect in our future research.

For the sake of comparison, we report SOTA performance in Table 2. Results from this research are in italics.

RQ4: Figure 3 shows that, on average, higher performance is associated with lower temperatures for both TempoWiC and HistoWiC, with accuracy decreasing as temperature values increase. Thus, we argue that high temperatures do not make it easier for GPT to solve WiC tasks or identify semantic changes effectively.

RQ5: ChatGPT Web results are presented in Table 3, along with the average performance we obtained through the GPT API across temperature values ranging from 0.0 to 1.0 (API 0–1), from 1.0 to 2.0 (API 1–2), and from 0.0 to 2.0 (API 0–2). As with GPT API, the performance of ChatGPT Web is higher for HistoWiC than for TempoWiC. In addition, our evaluation indicates that ChatGPT Web employs a moderate temperature setting, for we obtained consistent results when using a moderate temperature setting through GPT API. This suggests that the GPT API should be preferred for solving downstream task like WiC. It also suggests that the current SOTA evaluations may achieve higher results if the official API were used instead of the web interface. Thus, this implies that previous results using web interface should be interpreted with caution.

RQ6: As shown in Figure 4, the performance of (Chat)GPT Web decreases as the number of example messages increases (from ZSp to MSp).

Table 3: Comparison of GPT API and ChatGPT Web performance (Macro-F1)

	TempoWiC				HistoWiC			
	API 0–1	API 1–2	API 0–2	web	API 0–1	API 1–2	API 0–2	web
ZSp	.609	.589	.600	.580	.713	.665	.688	.686
FSp	.636	.606	.622	.569	.693	.626	.657	.674
MSp	-	-	-	.500	-	-	-	.565
all	.622	.598	.611	.550	.703	.645	.672	.642




Figure 5: Comparison of BERT Performance (Macro-F1) for TempoWiC and HistoWiC tasks across layers

This suggests that improving the performance of (Chat)GPT requires a more complex training approach than simply providing a few input-output examples. Furthermore, it indicates that the influence of message history is extremely significant in shaping the quality of conversations with (Chat)GPT. Indeed, a limited message history proved to be beneficial for the evaluation of GPT API through FSp.

RQ7: Figure 5 shows Macro-F1 scores obtained on TempoWiC and HistoWiC over the 12 BERT layers (see Appendix E). When considering the final layer, which is conventionally used in downstream tasks, BERT obtains Macro-F1 scores of .750 and .743 for TempoWiC and HistoWiC, respectively. Similar to Periti and Dubossarsky (2023), BERT performs best on HistoWiC when embeddings extracted from middle layers are considered. However, BERT performs best on TempoWiC when embeddings extracted from the last layers are used.

We compared the performance of GPT and BERT across their respective worst to best scenarios by sorting the Macro-F1 scores obtained by BERT and GPT in ascending order (bottom x-axis). For ChatGPT, we consider the results obtained through FSp and ZSp prompting for TempoWiC and HistoWiC, respectively. As shown in Figure 6, even when considering the best setting, GPT does not outperform the Macro-F1 score obtained by using the last layer of BERT, marked




Figure 6: GPT v BERT (Macro-F1). Performance is sorted in ascending order regardless of temperatures and layers. A black circle denotes the use of the last layer of BERT.




Figure 7: True Negative Rate v False Negative Rate. Each cross represents a (Chat)GPT experiment. Each dot represents the use of a specific layer of BERT.

with a black circle. However, although it exhibits lower performance, the results obtained from GPT are still comparable to BERT results on HistoWiC when embeddings extracted from the last layer of BERT are used.

Since our goal is to evaluate the potential of (Chat)GPT for recognizing lexical semantic changes, we analyzed the true negative rate and false negative rate scores, because *negative* examples represent semantic change in TempoWiC and HistoWiC datasets. As shown in Figure 7, regardless of the temperature and layer considered, (Chat)GPT falls short in recognizing semantic change for both TempoWiC and HistoWiC compared to BERT. However, it produces fewer false negatives than BERT for TempoWiC.

RQ8: In our experiment, GPT achieved low Spearman’s correlation coefficients for each temperature when ranking the target word of the LSC

English benchmark by degree of lexical semantic change. Higher correlations were achieved by using low temperatures rather than high ones (see Appendix F). Table 4 shows the GPT correlation for the temperature 0. For comparison, we report correlations obtained by BERT-based systems that leverage pretrained models. Note that, when BERT is fine-tuned, it generally achieves even higher correlation scores (see survey by Montanelli and Periti, 2023).

Table 4: LSC comparison: correlation obtained by SOTA, *pre-trained* BERT systems and GPT (temperature=0).

	Spearman’s correlation
Periti et al., 2023	.651
Laicher et al., 2021	.573
Periti et al., 2022	.512
Rother et al., 2020	.512
GPT API	.251

As shown in Table 4, the BERT-based system largely outperforms GPT, suggesting that GPT is not currently well-adapted for use in solving LSC downstream tasks.

5.1 BERT for Semantic Change Detection

There are notable differences between the Macro-F1 for TempoWiC and HistoWiC in terms of how the results increase and decrease across layers (see Figure 5). For TempoWiC the results increase until the 8th layer, after which they remain almost stable. Conversely, for HistoWiC the BERT performance rapidly increases until the 5th layer, after which it linearly decreases until the 12th layer. As regards Tempo WiC, we hypothesize that BERT is already aware of the set of word meanings considered for evaluation as it was pretrained on modern and contemporary texts. As regards HistoWiC, we hypothesize that BERT is not completely aware of the set of word meanings considered for evaluation and that word representations adopted for the historical context of HistoWiC¹² might be slightly tuned. Thus, using medium embedding layers could prove beneficial in detecting semantic changes, as these layers are less affected by contextualization (Ethayarajh, 2019). In other words, for HistoWiC, we hypothesize that the performance diminishes in the later layers due to the increasing contextualization of the medium and final embedding layers, which reduces the presence of noise in untuned word representations. This prompts us to question the ap-

¹²1810–1860, as referenced in Schlechtweg et al. (2020)

propriateness of using the last embedding layers to recognize historical lexical semantic change. We will address this question in future research.

6 Conclusion

In this study, we empirically investigated the use of the *current* (Chat)GPT 3.5 to detect semantic change. Our goal is not to comprehensively evaluate (Chat)GPT in dealing with semantic change, but rather to acknowledge its potential while also raising concerns and questions about its off-the-shelf use. In this regard, we used *reasonable* prompts from a human point of view, which may not necessarily be optimized for the model. We used the TempoWiC benchmark to assess (Chat)GPT’s ability to detect short-term semantic changes, and introduced a novel benchmark, HistoWiC, to assess (Chat)GPT’s ability to recognize long-term changes. When considering the standard 12 layer of BERT, our experiments show that (Chat)GPT achieves comparable performance to BERT (although slightly lower) in regard to detecting long-term changes, but performs significantly worse in regard to recognizing short-term changes. We find that BERT’s contextualized embeddings consistently provide a more effective and robust solution for capturing both short- and long-term changes in word meanings.

There are two possible explanations for the discrepancy in (Chat)GPT’s performance between TempoWiC and HistoWiC: i) HistoWiC might involve word meanings not explicitly covered during training, potentially aiding (Chat)GPT in detecting anomalies; ii) TempoWiC involves patterns typical of Twitter (now X), such as abbreviations, mentions, or tags, which may render it more challenging than HistoWiC.

In light of our findings, we argue that *(Chat)GPT 3.5 might be the hero the world deserves but not the one it needs right now*¹³, in particular for computationally modeling meaning over time, and by extension, for the study of semantic change. Nevertheless, during the course of our research, updates to (Chat)GPT became available and gained popularity, leading research and practitioners to conduct new experiments on these updated models. Particularly noteworthy is a recent study by Karjus (2023),

¹³This quote draws inspiration by the movie *Batman: The Dark Knight*. We leverage the analogy of (Chat)GPT achieving lower results than BERT to acknowledge the potential of (Chat)GPT while also raising concerns and questions about its use for Semantic Change detection.

which showcased remarkable performance on LSC using the GPT-4 model. Inspired by this research, our ongoing and future work is focused on further exploring the capabilities of GPT-4 for modeling semantic change.

Limitations

There are limitations we had to consider in the making of this paper. Firstly, a limitation arises when working with temporal HistoWiC benchmarks. While we ensure the utilization of diachronic data, we cannot guarantee that if the meaning of a word differs across contexts, it unequivocally indicates either the presence of stable polysemy (existing stable multiple meanings) or exemplifies a semantic change (either a new sense that it did not previously possess or a lost sense that it no longer has).

Other limitations are about the use of language models. We could not evaluate (Chat)GPT across different languages due to both price and API limitations. This means that while the results holds for English, we do not know how (Chat)GPT will behave for the other languages. Although we are aware of open source solution such as LLaMA, it still necessitates expensive research infrastructure, and we thus chose to focus on (Chat)GPT.

Like all research on (Chat)GPT (Laskar et al., 2023; Kocoń et al., 2023; Zhong et al., 2023), our work has a significant limitation that we cannot address: our (Chat)GPT results are not entirely reproducible as (Chat)GPT is inherently nondeterministic. In addition, like Zhong et al. (2023); Jiao et al. (2023), we found that time and economic constraints when using (Chat)GPT dictated that our evaluation of the software had to be based on only a subset of the TempoWiC and HistoWiC dataset.

In our study, we utilized (Chat)GPT 3.5. This could be considered a limitation, given the recent release of GPT 4. However, we opted for (Chat)GPT 3.5 based on the guidance provided in the current OpenAI documentation.¹⁴ Additionally, we argue that (Chat)GPT-3.5 is a cheaper alternative than the current GPT-4 model, making the investigation of (Chat)GPT-3.5 still significant for researchers with limited economic resources. We acknowledge that OpenAI continues to train and release new models, which could potentially affect the reproducibility of our results.

¹⁴<https://platform.openai.com/docs/guides/gpt/which-model-should-i-use>

One of the many features of (Chat)GPT is its ability to incorporate the history of preceding messages within a conversation while responding to new input prompts. However, there remain several unanswered questions regarding how this history influences the model’s answers. This holds true even for the zero-shot prompting strategy, where a general setting is lacking. Multiple prompts can be provided as part of the same chat or across different chats. For simplicity, and similar to previous research, we assigned only one chat for each ZSp experiment. We intend to use different chats in our future work to examine and investigate the effect of the message history.

Finally, as highlighted by Laskar et al. (2023), since the instruction-tuning datasets of OpenAI models are unknown (that is, not open source), the datasets used for evaluation may or may not be part of the instruction-tuning training data of OpenAI.

Despite these limitations, we argue that our work is significant as it may prompt new discussion on the use of LMs such as BERT and (Chat)GPT, while also dispelling the expanding belief that the use of ChatGPT as *off-the-shelf* model *already* makes BERT an outdated technology.

Acknowledgements

This work has in part been funded by the project Towards Computational Lexical Semantic Change Detection supported by the Swedish Research Council (2019–2022; contract 2018-01184), and in part by the research program Change is Key! supported by Riksbankens Jubileumsfond (under reference number M21-0021).

References

- Ian L. Alberts, Lorenzo Mercolli, Thomas Pyka, George Prenosil, Kuangyu Shi, Axel Rominger, and Ali Afshar-Oromieh. 2023. [Large Language Models \(LLM\) and ChatGPT: What Will the Impact on Nuclear Medicine Be?](#) *European journal of nuclear medicine and molecular imaging*, 50(6):1549–1552.
- Fares Antaki, Samir Touma, Daniel Milad, Jonathan El-Khoury, and Renaud Duval. 2023. [Evaluating the Performance of ChatGPT in Ophthalmology: An Analysis of Its Successes and Shortcomings.](#) *Ophthalmology Science*, 3(4):100324.
- Nikolay Arefyev, Maksim Fedoseev, Vitaly Probstov, Daniil Homiskiy, Adis Davletov, and Alexander Panchenko. 2021. [DeepMistake: Which Senses are Hard to Distinguish for a Word-in-Context Model.](#) In *Proceedings of the Conference on Computational Linguistics and Intellectual Technologies (Dialogue)*, (online). RSUH.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. [A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity.](#)
- Leonard Bloomfield. 1933. *Language*. New York: Holt, Rinehart & Winston.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners.](#) In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Pierluigi Cassotti, Lucia Siciliani, Marco DeGemmis, Giovanni Semeraro, and Pierpaolo Basile. 2023. [XL-LEXEME: WiC Pretrained Model for Cross-Lingual LEXical sEMantic change.](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1577–1585, Toronto, Canada. Association for Computational Linguistics.
- Ze Chen, Kangxu Wang, Zijian Cai, Jiewen Zheng, Jiarong He, Max Gao, and Jason Zhang. 2022. [Using Deep Mixture-of-Experts to Detect Word Meaning Shift for TempoWiC.](#)
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language](#)

- [Understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kawin Ethayarajh. 2019. [How Contextual are Contextualized Word Representations? Comparing the Geometry of BERT, ELMo, and GPT-2 Embeddings](#). 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 55–65, Hong Kong, China. Association for Computational Linguistics.
- Dirk Geeraerts, Dirk Speelman, Kris Heylen, Mariana Montes, Stefano De Pascale, Karlien Franco, and Michael Lang. 2024. [Lexical Variation and Change: A Distributional Semantic Approach](#). Oxford University Press.
- A. Shaji George and A. S. Hovan George. 2023. [A Review of ChatGPT AI’s Impact on Several Business Sectors](#). *Partners Universal International Innovation Journal*, 1(1):9–23.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. [Is ChatGPT A Good Translator? Yes With GPT-4 As The Engine](#).
- Andres Karjus. 2023. [Machine-assisted Mixed Methods: Augmenting Humanities and Social Sciences with Artificial Intelligence](#).
- Margarita Kay. 1979. [Lexemic Change and Semantic Shift in Disease Names](#). *Culture, Medicine and Psychiatry*, 3(1):73–94.
- Mohammad Khalil and Erkan Er. 2023. [Will chatgpt get you caught? rethinking of plagiarism detection](#). In *Learning and Collaboration Technologies: 10th International Conference, LCT 2023, Held as Part of the 25th HCI International Conference, HCII 2023, Copenhagen, Denmark, July 23–28, 2023, Proceedings, Part I*, page 475–487, Copenhagen, Denmark. Springer-Verlag.
- Jan Kocoń, Igor Cichecki, Oliwier Kaszyca, Mateusz Kochanek, Dominika Szydło, Joanna Baran, Julita Bielaniewicz, Marcin Gruza, Arkadiusz Janz, Kamil Kanclerz, Anna Kocoń, Bartłomiej Koptyra, Wiktoria Mieleszczenko-Kowszewicz, Piotr Miłkowski, Marcin Oleksy, Maciej Piasecki, Łukasz Radliński, Konrad Wojtasik, Stanisław Woźniak, and Przemysław Kazienko. 2023. [ChatGPT: Jack of All Trades, Master of None](#). *Information Fusion*, 99:101861.
- Andrey Kutuzov, Lilja Øvrelid, Terrence Szyman-ski, and Erik Velldal. 2018. [Diachronic Word Embeddings and Semantic Shifts: a Survey](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1384–1397, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Viet Lai, Nghia Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Nguyen. 2023. [ChatGPT Beyond English: Towards a Comprehensive Evaluation of Large Language Models in Multilingual Learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13171–13189, Singapore. Association for Computational Linguistics.
- Severin Laicher, Sinan Kurtyigit, Dominik Schlechtweg, Jonas Kuhn, and Sabine Schulte im Walde. 2021. [Explaining and Improving BERT Performance on Lexical Semantic Change Detection](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 192–202, Online. Association for Computational Linguistics.
- Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Huang. 2023. [A Systematic Study and Comprehensive Evaluation of ChatGPT on Benchmark Datasets](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 431–469, Toronto, Canada. Association for Computational Linguistics.
- Meng Liang and Yao Shi. 2023. [Named Entity Recognition Method Based on BERT-whitening and Dynamic Fusion Model](#). In *2023 5th International Conference on Natural Language Processing (ICNLP)*, pages 191–197.


- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. [Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation](#).
- Daniel Loureiro, Aminette D’Souza, Areej Nasser Muhajab, Isabella A. White, Gabriel Wong, Luis Espinosa-Anke, Leonardo Neves, Francesco Barbieri, and Jose Camacho-Collados. 2022. [TempoWiC: An Evaluation Benchmark for Detecting Meaning Shift in Social Media](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3353–3359, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Brady D Lund and Ting Wang. 2023. [Chatting about ChatGPT: How May AI and GPT Impact Academia and Libraries?](#) *Library Hi Tech News*, 40(3):26–29.
- Chenyang Lyu, Yongxin Zhou, and Tianbo Ji. 2022. [MLLabs-LIG at TempoWiC 2022: A Generative Approach for Examining Temporal Meaning Shift](#). In *Proceedings of the The First Workshop on Ever Evolving NLP (EvoNLP)*, pages 1–6, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Xiaofei Ma, Zhiguo Wang, Patrick Ng, Ramesh Nallapati, and Bing Xiang. 2019. [Universal Text Representation from BERT: An Empirical Study](#).
- Stefano Menini, Teresa Paccosi, Sara Tonelli, Marieke Van Erp, Inger Leemans, Pasquale Lisena, Raphael Troncy, William Tullett, Ali Hürriyetoglu, Ger Dijkstra, Femke Gordijn, Elias Jürgens, Josephine Koopman, Aron Ouwerkerk, Sanne Steen, Inna Novalija, Janez Brank, Dunja Mladenec, and Anja Zidar. 2022. [A Multilingual Benchmark to Capture Olfactory Situations over Time](#). In *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*, pages 1–10, Dublin, Ireland. Association for Computational Linguistics.
- Stefano Montanelli and Francesco Periti. 2023. [A Survey on Contextualised Semantic Shift Detection](#).
- Teresa Paccosi, Stefano Menini, Elisa Leonardelli, Ilaria Barzon, and Sara Tonelli. 2023. [Scent and Sensibility: Perception Shifts in the Olfactory Domain](#). In *Proceedings of the 4th Workshop on Computational Approaches to Historical Language Change*, pages 143–152, Singapore. Association for Computational Linguistics.
- Keqin Peng, Liang Ding, Qihuang Zhong, Li Shen, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng Tao. 2023. [Towards Making the Most of ChatGPT for Machine Translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5622–5633, Singapore. Association for Computational Linguistics.
- Francesco Periti and Haim Dubossarsky. 2023. [The Time-Embedding Travelers@WiC-ITA](#). In *Proceedings of the Eighth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2023)*, Parma, Italy. CEUR.org.
- Francesco Periti, Alfio Ferrara, Stefano Montanelli, and Martin Ruskov. 2022. [What is Done is Done: an Incremental Approach to Semantic Shift Detection](#). In *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*, pages 33–43, Dublin, Ireland. Association for Computational Linguistics.
- Francesco Periti, Sergio Picascia, Stefano Montanelli, Alfio Ferrara, and Nina Tahmasebi. 2023. [Studying Word Meaning Evolution through Incremental Semantic Shift Detection: A Case Study of Italian Parliamentary Speeches](#).
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. [WiC: the Word-in-Context Dataset for Evaluating Context-Sensitive Meaning Representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B Viegas, Andy Coenen, Adam Pearce, and Been Kim. 2019. [Visualizing and Measuring the Geometry of BERT](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

- David Rother, Thomas Haider, and Steffen Eger. 2020. [CMCE at SemEval-2020 Task 1: Clustering on Manifolds of Contextualized Embeddings to Detect Historical Meaning Shifts](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 187–193, Barcelona (online). International Committee for Computational Linguistics.
- Dominik Schlechtweg, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky, and Nina Tahmasebi. 2020. [SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1–23, Barcelona (online). International Committee for Computational Linguistics.
- Dominik Schlechtweg, Sabine Schulte im Walde, and Stefanie Eckmann. 2018. [Diachronic Usage Relatedness \(DUREl\): A Framework for the Annotation of Lexical Semantic Change](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 169–174, New Orleans, Louisiana. Association for Computational Linguistics.
- Dominik Schlechtweg, Nina Tahmasebi, Simon Hengchen, Haim Dubossarsky, and Barbara McGillivray. 2021. [DWUG: A large Resource of Diachronic Word Usage Graphs in Four Languages](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7079–7091, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Dominik Schlechtweg, Shafqat Mumtaz Virk, Pauline Sander, Emma Sköldborg, Lukas Theuer Linke, Tuo Zhang, Nina Tahmasebi, Jonas Kuhn, and Sabine Schulte im Walde. 2023. [The DUREl Annotation Tool: Human and Computational Measurement of Semantic Proximity, Sense Clusters and Semantic Change](#).
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. [HuggingGPT: Solving AI Tasks with ChatGPT and its Friends in Hugging Face](#).
- Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2021. [Survey of computational approaches to lexical semantic change detection](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Leonie Weissweiler, Valentin Hofmann, Anjali Kantharuban, Anna Cai, Ritam Dutt, Amey Hengle, Anubha Kabra, Atharva Kulkarni, Abhishek Vijayakumar, Haofei Yu, Hinrich Schuetze, Kemal Oflazer, and David Mortensen. 2023. [Counting the Bugs in ChatGPT’s Wugs: A Multilingual Investigation into the Morphological Capabilities of a Large Language Model](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6508–6524, Singapore. Association for Computational Linguistics.
- Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du, and Dacheng Tao. 2023. [Can ChatGPT Understand Too? A Comparative Study on ChatGPT and Fine-tuned BERT](#).
- Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, Hao Peng, Jianxin Li, Jia Wu, Ziwei Liu, Pengtao Xie, Caiming Xiong, Jian Pei, Philip S. Yu, and Lichao Sun. 2023. [A Comprehensive Survey on Pretrained Foundation Models: A History from BERT to ChatGPT](#).

Appendix

A Historical WiC

We shifted from the LSC to the WiC setting as follows. First, we selected only the annotated LSC instances containing contexts from different time periods. We then filtered out all the instances annotated by a single annotator¹⁵ and all the instances that are associated with an average score, s , such that $1.5 < s < 3.5$, which represents ambiguous cases even for humans. Finally, we binarized the LSC annotations by converting each $s \leq 1.5$ to *False* (i.e. different meanings) and each $s \geq 3.5$ to *True* (i.e. same meaning). We report in Table 5 the four-point semantic-relatedness used to annotate the LSC instances through the DUREl framework.



4:	Identical
3:	Closely related
2:	Distantly related
1:	Unrelated

Table 5: The DUREl relatedness scale used in [Schlechtweg et al. \(2020, 2018\)](#)

B Message history

Although one of the many features of (Chat)GPT is its ability to consider the history of preceding messages within a conversation while responding to new input prompts, GPT API and the web version handle message history differently. In GPT API, the message history is limited to a fixed number of tokens (i.e., 4,096 tokens for gpt-3.5-turbo); however, we are not aware of how the message history is handled in ChatGPT Web, where an unlimited number of message for chat seems to be supported.

In our experiments, we use a single chat for each considered prompting strategy, both for ChatGPT Web and GPT API. However, in ChatGPT Web, we considered the full message history for the ZSp, FSp, and MSp strategies. Instead, to avoid exceeding the token limit set by the OpenAI API, we tested GPT API for the ZSp and FSp strategies by considering a message history of 33 messages. Note that due to the token limit, testing the MSp strategy for GPT API wasn't possible, as the limited message history would make MSp equivalent to FSp. The 33-message history was organized as

a combination of a *fixed* and a *sliding window*. We set the fixed window to ensure the model is always aware of the task we asked it to answer in the early prompts; instead, we set the sliding window to emulate the flow of the conversation as in ChatGPT Web. In particular, i) in ZSp, the fixed window covers our first prompt (i.e., task explanation) and the (Chat)GPT answer, while the sliding window covers the i -th prompts and the last 30 messages (i.e., 15 prompts and 15 (Chat)GPT answers); ii) in FSp, the fixed window covers the first 26 messages (i.e., task explanation and example instances), while the sliding window covers the i -th prompts and the last 6 messages. Figure 8 summarizes the message history we set for testing GPT API.

¹⁵Different instances were annotated by varying numbers of annotators.

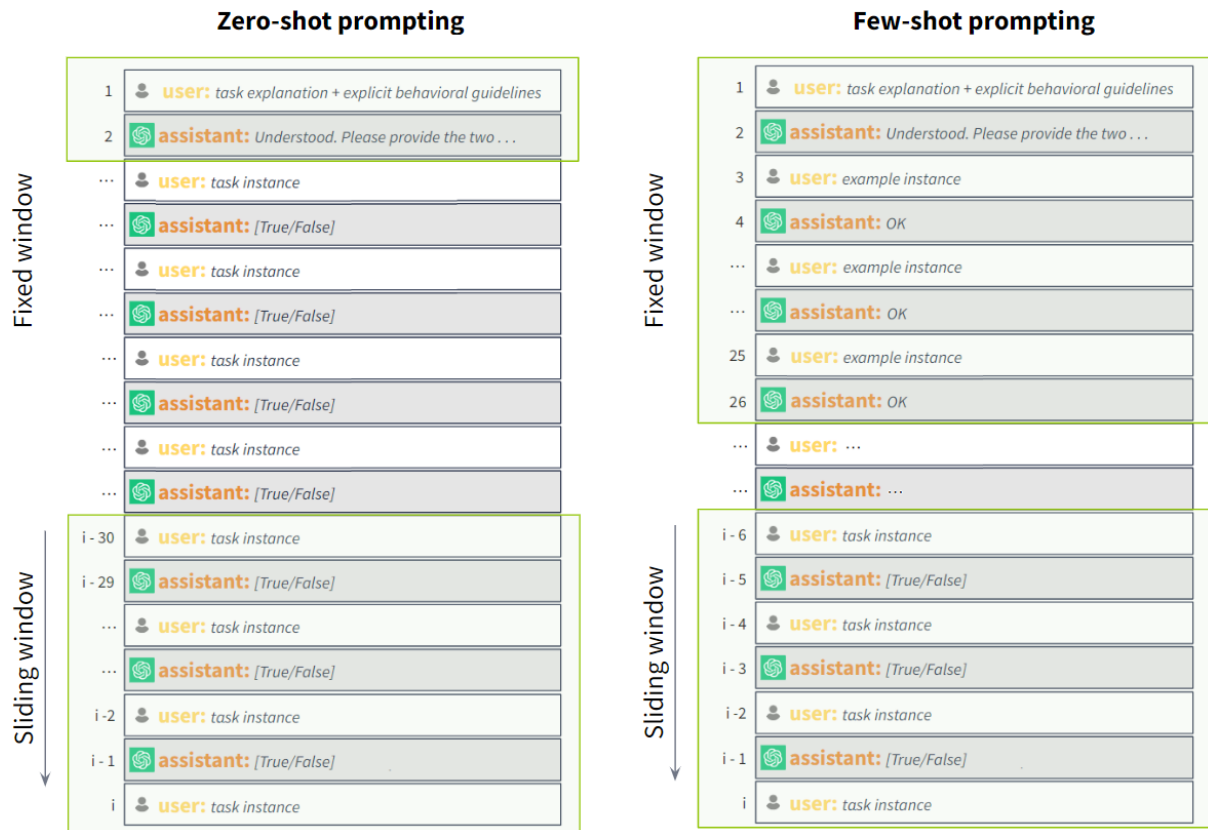


Figure 8: Message history used for GPT API in the zero-shot prompting (ZSp) and few-shot prompting (FSp) strategies. The message history is organized as a combination of a fixed and a sliding window, encompassing a total of 33 messages. The fixed window ensures that the model remains constantly aware of the task we have asked it to address in the initial prompts and the given examples (if any). Conversely, we establish the sliding window to emulate the conversational flow of ChatGPT Web.

C (Chat)GPT templates

C.1 WiC template

Description	Template
task explanation	Task: Determine whether two given sentences use a target word with the same meaning or different meanings in their respective contexts.
explicit behavioral guidelines	I'll provide some negative and positive examples to teach you how to deal with the task before testing you. Please respond with only "OK" during the examples; when it's your turn, answer only with "True" or "False" without any additional text. When it's your turn, choose one: "True" if the target word has the same meaning in both sentences; "False" if the target word has different meanings in the sentences. I'll notify you when it's your turn.
example instance	This is an example. You have to answer "OK": Sentence 1: [First sentence containing the target word] Sentence 1: [First sentence containing the target word] Target: [Target word] Question: Do the target word in both sentences have the same meaning in their respective contexts? Answer: [True/False]
task instance	Now it's your turn. You have to answer with "True" or "False": Sentence 1: [First sentence containing the target word] Sentence 1: [First sentence containing the target word] Target: [Target word] Question: Do the target word in both sentences have the same meaning in their respective contexts? Answer: [The model is expected to respond with "True" or "False"]

Table 6: Sections of the prompt template used for testing (Chat)GPT.

ID	Strategy	Prompt
ZSp	zero-shot prompting	task explanation explicit behavioral guidelines task instance ... task instance
FSp	few-shot prompting	task explanation explicit behavioral guidelines example instance ... example instance task instance ... task instance
MSp	many-shot prompting	<i>like FSp</i>

Table 7: Prompt template for each employed prompting strategy.

C.2 LSC template

Strategy	Template
ZSp	Consider the following two time periods and target word. How much has the meaning of the target word changed between the two periods? Rate the lexical semantic change on a scale from 0 to 1. Provide only a score. Target: [Target word] Time period 1: 1810–1860 Time period 2: 1960–2010 Answer: [The model is expected to respond with a continuous score s , with $0 \leq s \leq 1$]

Table 8: Prompt template for LSC.

D GPT API performance on TempoWiC and HistoWiC

D.1 Experiment 1 - temperature

		GPT API - Temperature												
		prompt	0.0	0.2	0.4	0.6	0.8	1.0	1.2	1.4	1.6	1.8	2.0	avg
TempoWiC	ZSp		.568	.584	.604	.599	.592	.576	.604	.560	.560	.599	.579	.584
	FSp		.648	.648	.664	.634	.597	.631	.645	.585	.608	.581	.598	.622
HistoWiC	ZSp		.728	.683	.689	.676	.666	.694	.715	.609	.704	.671	.594	.675
	FSp		.684	.698	.721	.698	.671	.700	.686	.599	.552	.607	.601	.656

Table 9: Comparison of GPT performance (Macro-F1) for TempoWiC and HistoWiC at various temperature values using the official API and different prompts.

D.2 Experiment 2 - temperature

		GPT API - Temperature												
		prompt	0.0	0.2	0.4	0.6	0.8	1.0	1.2	1.4	1.6	1.8	2.0	avg
TempoWiC	ZSp		.645	.628	.643	.605	.664	.602	.600	.598	.575	.580	.636	.616
	FSp		.659	.632	.649	.627	.644	.597	.689	.627	.597	.551	.562	.621
HistoWiC	ZSp		.751	.758	.711	.765	.729	.712	.678	.652	.679	.664	.604	.700
	FSp		.684	.678	.707	.700	.706	.665	.607	.662	.615	.592	.623	.658

Table 10: Comparison of GPT performance (Macro-F1) for TempoWiC and HistoWiC at various temperature values using the official API and different prompts.

D.3 Average performance per temperature

		GPT API - Temperature												
prompt		0.0	0.2	0.4	0.6	0.8	1.0	1.2	1.4	1.6	1.8	2.0	avg	
TempoWiC	ZSp	.606	.606	.624	.602	.628	.589	.602	.579	.568	.589	.607	.600	
	FSp	.654	.640	.657	.631	.620	.614	.667	.606	.602	.566	.580	.622	
HistoWiC	ZSp	.740	.720	.700	.720	.698	.703	.696	.631	.692	.668	.599	.688	
	FSp	.684	.688	.714	.699	.688	.682	.647	.631	.584	.599	.612	.657	

Table 11: Comparison of GPT performance (Macro-F1) for TempoWiC and HistoWiC at various temperature values using the official API and different prompts. We report the average performance for each temperature.

E BERT performance on TempoWiC and HistoWiC

		Layers												avg
		1	2	3	4	5	6	7	8	9	10	11	12	
TempoWiC		.669	.631	.635	.627	.604	.627	.704	.749	.744	.730	.737	.751	.684
HistoWiC		.650	.678	.739	.782	.828	.801	.806	.771	.771	.749	.722	.744	.753

Table 12: Comparison of BERT Performance (Macro-F1) for TempoWiC and HistoWiC tasks at different embedding layers.

F GPT API performance on LSC

		Temperature											avg
		0.0	0.2	0.4	0.6	0.8	1.0	1.2	1.4	1.6	1.8	2.0	
SemEval-English		.251	.200	.207	.279	.008	.012	.230	.154	.011	.194	.004	

Table 13: Comparison of (Chat)GPT performance (Spearman’s correlation) for LSC on SemEval-English at various temperature values using the official API.