



HAL
open science

Prompt Selection Matters: Enhancing Text Annotations for Social Sciences with Large Language Models

Louis Abraham, Charles Arnal, Antoine Marie

► **To cite this version:**

Louis Abraham, Charles Arnal, Antoine Marie. Prompt Selection Matters: Enhancing Text Annotations for Social Sciences with Large Language Models. 2024. hal-04657856

HAL Id: hal-04657856

<https://hal.science/hal-04657856>

Preprint submitted on 22 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Prompt Selection Matters: Enhancing Text Annotations for Social Sciences with Large Language Models

Louis Abraham^{*1}, Charles Arnal^{*2}, and Antoine Marie^{*3}

¹Université Paris 1 Panthéon-Sorbonne

²Université Paris-Saclay, CNRS, Inria

³Institut Jean Nicod, Ecole Normale Supérieure, PSL-EHESS-CNRS

Abstract

Large Language Models have recently been applied to text annotation tasks from social sciences, equalling or surpassing the performance of human workers at a fraction of the cost. However, no inquiry has yet been made on the impact of prompt selection on labelling accuracy. In this study, we show that performance greatly varies between prompts, and we apply the method of automatic prompt optimization to systematically craft high quality prompts. We also provide the community with a simple, browser-based implementation of the method at <https://prompt-ultra.github.io/>.

1 Introduction

Throughout the social sciences, many research questions are answered through annotation and classification of large volumes of text, such as tweets or Facebook comments. Researchers may be interested in knowing, for instance, how politically slanted (liberal vs. conservative) a claim or headline is; how emotional or hostile its tone is, or which one of the basic emotions it reflects (sadness, joy, anger, etc) [BLHJ16, GL12, SM10, RBOP24]. Text annotation so far had to be performed either by human experts or by unskilled crowd workers, depending on the nature of the task. As a result, it was typically costly and time-consuming; crowd workers are also likely to mislabel the data.

The recent progress of Large Language Models (LLMs) has opened new avenues and could revolutionize text mining by allowing huge volumes of text data to be analyzed in an unsupervised way in a matter of minutes [HCvH24, GAK23, Tör23, WR23]. Early results suggest that LLMs can perform extremely well, extremely fast and at almost no financial cost, with levels of accuracy rivalling those of experts and superior to those of unskilled workers. Moreover, these results are achieved by off-the-shelf, general purpose models, that do not need any specialized training, unlike some pre-existing machine learning-based techniques.

Nonetheless, crucial aspects of the automatic annotation of text data using LLMs have not been studied yet. In particular, earlier studies did not consider the importance of prompt choice: they used simple hand-crafted prompts, such as “Does the following message express liberal or conservative views?” or “Is this tweet pro-life or pro-choice?”, to annotate their corpus. However, it has been observed outside the context of text annotation that applying distinct yet similar versions of a given prompt to certain tasks can result in large differences in accuracy, of the order of more than 10% [KGR⁺22, BMR⁺20]. Going from 10% to 20% of mislabelled data can greatly impact the quality of a study’s conclusions, especially if the classification errors are biased (e.g., if all mislabelled tweets express conservative views).

^{*}Equal contribution.

Our main objectives in this paper are threefold. First, to help social scientists understand and adopt automatic text annotation using LLMs by providing a clear illustration of it. Second, to investigate the importance of prompt selection on performance, and to raise awareness about the issue of performance variability. Third, to explain to social scientists how to apply state-of-the-art prompt optimization methods used in the wider LLM community [ZMH⁺22, YWL⁺24] to their own classification tasks.

More precisely, our contributions are as follows:

- We give a short, didactic overview of automatic text annotation for social sciences using LLMs; in particular, we draw attention to some shortcomings of the method that have not yet been discussed.
- We describe the principle of automatic prompt optimization, and provide a simple implementation of the method.
- We investigate and quantify the impact of prompt selection on accuracy across a range of standard classification tasks in social sciences using LLMs. To that end, we compare both hand-crafted prompts and automatically optimized prompts.
- We conclude that apparently similar prompts yield greatly varied accuracy levels. We also observe that automatic prompt optimization yields consistently good performance and beats prompt-crafting heuristics (such as using Chain of Thoughts prompts, see below) on most tasks.
- Finally, we provide the community with a simple and efficient way to label their datasets using LLMs coupled with either their own hand-crafted prompts or with a prompt optimization algorithm: our browser-based service <https://prompt-ultra.github.io/>.

All of the code used will be made available on GitHub.

2 Automatic text labelling using LLMs

In this section, we present the recently proposed method of automatic text labelling using LLMs, and discuss some of its shortcomings.

State-of-the-art general purpose language models, such as GPT-4 [Ope23], have now achieved human-like performance on a variety of tasks. In particular, it has been recently demonstrated that they can be applied off-the-shelf, i.e., without needing any specialized training, to a variety of text-labelling tasks needed in social sciences. Among other examples, they have been shown to reach around 95% of accuracy¹ on negativity detection in tweets and news articles [HCvH24], around 75% of accuracy on stance detection in tweets (largely outperforming untrained crowd workers) [GAK23], and more than 92% of accuracy on political orientation in tweets [Tör23].

Up until one or two years ago, similar performances were either unachievable, or only achievable by models specifically trained on the task at hand, the creation of which was complex and time-consuming [JD16, NR13, AAS⁺20, PP11, DMB23].

Let us give a quick explanation of how to automatically label text using LLMs. Though variants are possible, the simplest method is to simply ask in simple and clear terms an LLM-based chatbot, such as ChatGPT [LHM⁺23], to give a label to a given piece of text. For instance:

```
[Social scientist]
Does the following tweet express liberal or conservative opinions?
Output only "liberal" or "conservative" without quotes.
```

```
Please support our Capitol Police and Law Enforcement.
They are truly on the side of our Country. Stay peaceful!
```

```
[ChatGPT] conservative
```

¹It is useful to note in passing that while it is desirable to maximise accuracy at classification tasks, it is generally impossible to reach 100% in virtue of the fact that some statements are intrinsically ambiguous and thus impossible to classify even by professional raters.

Typing each prompt by hand in the browser interface would be tedious. Instead, researchers would typically send such prompts through the API using e.g., a simple Python script², allowing thousands of statements to be classified in a matter of minutes. Each message would have to be processed in a fresh instance of the client chat, to avoid a progressive biasing of the model by the conversation history. A small percentage of answers might not be correctly formatted (the LLM might respond “The message is conservative” or “Conservative.” or “right-wing” instead of “conservative”), in which case one would simply have to ask the chatbot again. To get an estimate of the accuracy of the labels thus produced, the researcher should ideally have a small subset of the dataset be annotated by experts (typically a few hundreds to a few thousand items of text), so that the labels provided by the LLM can be compared to those provided by experts, which serve as ground truth for the task. The degree of agreement between the two sets of labels can then serve as an estimate of the accuracy of the LLM’s annotations.

The perks of using LLMs to classify text, compared to manual annotation by crowd workers or experts, are immense. It is easy to implement, fast and inexpensive (thousands of short or medium length texts can be labelled in a few minutes for less than 10\$). It also outperforms earlier machine learning-based methods without requiring the training of a specialized model, which can be time-consuming and requires technical skills.

2.1 Limitations of automatic text labelling

As for any method, automatic text annotation using LLMs suffers from some limitations, some of which have not been extensively discussed in the literature. Though early studies are very encouraging, the range of tasks that have been studied so far remains rather limited; the vast majority of them consists in attributing one of a handful of labels to short- to medium-length texts (most of which being social media posts or news headlines). Moreover, some preliminary results find that accuracy tends to be lower in languages other than English [HCvH24], which is coherent with general observations regarding LLMs’ performances. LLMs are also known to suffer from certain biases, including political ones [LJW⁺22, MPNR24, ZLS⁺24]; those could result in systematic biases on certain tasks that would have a greater impact on the research performed on the annotated dataset than the raw accuracy might suggest. As an example, consider an LLM tasked with labelling a set of tweets as Republican-leaning or Democrat-leaning ; if all mislabelled tweets are Republican tweets erroneously labeled as Democrat, the conclusions drawn from the dataset might be wrong, despite the overall accuracy being high.

The issue can also be exacerbated by conscious moderating efforts by the LLM’s creators: various techniques are in place to ensure that the answers provided by LLM-based chat agents avoid using stereotypes susceptible to be seen as offensive, which could negatively affect their performance on some politically sensitive tasks, such as race.

Other problems arise not from the intrinsic nature of LLMs, but rather from the way in which they are currently being trained and deployed. State-of-the-art LLMs, such as GPT-4, are trained on immense datasets of text scrapped from the web, and regularly updated in a similar fashion. Beyond the possible privacy issues, this has crucial implications regarding the performance of LLMs: they will often perform better on data on which they have been trained than on yet unseen data (e.g., data that was produced after their training). As a result, the same LLM might be much better at classifying a certain set of tweets than a similar yet more recent set which it has not yet “seen”. As the exact nature of the training data is typically not shared with the general public, such phenomena are very hard to predict - see some of our observations in Section 4. In particular, we suspect that some of the most impressive accuracy scores achieved by LLMs on annotation tasks and reported in the literature might be partially explained by this. This problem, in conjunction with the regular updating of models, creates important issues of replicability: results might vary greatly between apparently similar tasks, or when applying different versions of the same model to the same task. This can be partially offset by using “frozen” LLMs, i.e. by keeping a copy of a certain version of an LLM at a certain point in time and using it on all tasks.

²Such a script can be found in our code base.

3 Experiments

We evaluate various hand-crafted prompts and automatically optimized prompts (see Subsections 3.2 and 3.3 below) on a range of datasets and classical social sciences tasks using OpenAI’s GPT-3.5 Turbo’s API [YCX⁺23] (see our code base for details). Our objective is twofold: first, we want to measure the impact of the precise formulation of the prompt on GPT’s accuracy by comparing performance across various prompts. The goal here is not to find some optimal prompt crafting technique that would systematically outperform all others (such panacea is unlikely to exist, as suggested by our experimental results further below); in particular, we do not claim to have examined every possible prompt-crafting trick. Rather, we only want to see whether two reasonable, semantically similar prompts can result in significantly different accuracies on a given task. Second, we want to check whether automatic prompt optimization can help achieve consistently good (though not necessarily optimal) results on all tasks without the need for manual tweaking by the experimenter.

3.1 Datasets and tasks

We have selected a range of diverse yet typical annotation tasks and datasets. When the sets have a predefined train and test sets split, we only use the test set for reasons explained in Section 4.

TweetEval - hate, emotion, sentiment, offensive (TE-hate, TE-emotion, TE-sent, TE-off)

TweetEval [BCCEAN20] consists of seven heterogeneous tasks performed on Twitter data in English, of which we have selected four:

- Hate detection (hateful or non-hateful); the set contains 2’970 tweets [BBF⁺19].
- Emotion recognition (anger, joy, optimism or sadness); the set contains 1’421 tweets [MBMSK18].
- Sentiment recognition (negative, neutral, positive); the set contains 12’284 tweets, and we randomly draw and use 10’000 of them [RFN17].
- Offensive language detection (non-offensive, offensive); the set contains 860 tweets [ZMN⁺19].

Tweet Sentiment Multilingual (TML-sent)

Tweet Sentiment Multilingual [BEACC22] is a dataset of tweets in 8 different languages (Arabic, English, French, German, Hindi, Italian, Portuguese, Spanish) with sentiment analysis labels (negative, neutral and positive). The test set contains 6’960 messages.

Article Bias Prediction (AS-pol)

A dataset of 37554 news articles from major US newspapers [BDSMGN20], with labels representing their political inclination (left, center³ or right), from which we randomly sample 10’000 articles.

Liberals vs Conservatives on Reddit (LibCon)

A dataset of 13’000 Reddit posts in English collected from liberal- and conservative-leaning subreddits [lib]. We randomly select and use 10’000 of them.

3.2 Hand-crafted prompts

It has been observed outside the context of automatic text annotation that various tricks can help increase LLMs’ performance. As an example, simply adding “Let’s think step by step” at the end of a prompt can result in impressive improvements [KGR⁺22]. Though there is no systematic theory of “good prompting”, many such tricks have been catalogued (see e.g., [pro]). We have applied several of these standard techniques to craft five different prompts for each task; below are examples from the **TweetEval - hate** task. The other prompts can be found in the accompanying code.

- **Simple** - A simple, minimalist prompt:
Classify the following message as hateful if it contains hate speech or non-hateful if it does not contain hate speech. Output only “hateful” or “non-hateful” without quotes.

³Note that the centrist label describes articles that are biased towards a centrist political ideology, and not articles that lack political bias.

- **Explanations** - A prompt enriched with additional explanations (see e.g. [LDC⁺22]):
You will be given a tweet. Your task is to use your knowledge of sociology, politics and culture to make an educated guess on whether the tweet contains hate speech. Hate speech includes any speech that expresses hate or encourages violence towards a person or group based on something such as race, religion, sex, sexual orientation or socio-economic situation. For example, messages that aggressively and unfairly target women or immigrants constitute hate speech. If the message does not contain enough information to correctly classify it, just make your best guess. If the message contains hate speech, output “hateful” without quotes. If the message does not contain hate speech, output “non-hateful” without quotes. Output only “hateful” or “non-hateful” without quotes.
- **Examples** - A prompt containing the same explanations, as well as some examples of correctly classified messages (see e.g. [BMR⁺20]).
You will be given a tweet. Your task is to [...]
Output only “hateful” or “non-hateful” without quotes.
Here are four examples of tweets with the correct classification :
Women are equal and deserve respect. Just kidding, they should suck my dick.
hateful
I’ve been called a son of a whore and a cunt today already. It’s not even 1 yet
non-hateful
[...]
- **Roleplay** - A prompt containing the same explanations and asking the chatbot to answer the question while roleplaying as a political analyst (see e.g. [BG24]).
You will be given a tweet. Pretend that you are a political analyst and decide whether the tweet contains hate speech. Hate speech includes [...]
- **Chain of thoughts (CoT)** - A prompt with the same explanations that asks the LLM to first explain its reasoning, then answer the question (see [WWS⁺22, KGR⁺22]). The very last word of the output is used as a label.
You will be given a tweet. Your task is to [...] constitute hate speech. First explain your reasoning. Then start a new line, and output “hateful” without quotes if the message contains hate speech, and output “non-hateful” without quotes if the message does not contain hate speech. The only word on the new line must be either “hateful” without quotes or “non-hateful” without quotes.

3.3 Automatic prompt optimization (APO)

A more systematic way to craft good prompts has been recently proposed: automatic prompt optimization [ZMH⁺22, YWL⁺24, SDZ23] (also called automatic prompt engineering).

The core idea is to ask an LLM to repeatedly rephrase prompts, then to select the one that offers the best performance. We translated this general concept into our specific text labelling setting as follows: we first require a subset of the dataset to have already been labelled by human annotators; typically, a few hundred or a few thousand messages. We then start from any reasonable prompt for the task, e.g., “Classify the following message as hateful if it contains hate speech or non-hateful if it does not contain hate speech.” in the case of the hate detection task. We ask the LLM to reformulate the prompt several times, e.g. with “Generate a variation of the following instruction while keeping the semantic meaning.” We then repeat the following steps as many times as desired: we evaluate the current set of prompts on the labelled subset, we keep only the top prompts (in terms of scores), and we ask the LLM to reformulate them to create a new generation of prompts (in addition with the top prompts themselves). The best prompt of the last generation is then used to label the remainder of the dataset. Before that, one can estimate its associated accuracy by testing it on another pre-labelled subset, distinct from the one used during the optimization process to avoid any bias caused by a lack of independence.

In our experiments and for a given task, each generation had 8 prompts; each was evaluated on a fixed subset of 400 samples from the dataset. The top 2 prompts were kept, and each was reformulated 3 times to generate a new generation of $2 + 2 \times 3 = 8$ prompts. We repeated this process for 15 generations, after which we kept the best performing prompt of the last generation. Finally, we

evaluated this prompt on the remainder of the dataset (without the 400 prompts, to guarantee an unbiased estimate of the prompt’s accuracy).

The method is both conceptually simple, and easy to implement; its main a priori downside compared to using hand-crafted prompts is the increased number of calls to the chatbot API required, though the costs remain almost negligible. It can be tested on our free browser-based service, which we further describe in Appendix A.

4 Results and discussion

Accuracies

	TE-hate	TE-emotion	TE-sent	TE-off	TML-sent	AS-pol	LibCon
Simple	<i>57.0</i>	78.7	<i>60.5</i>	<i>71.4</i>	67.1	52.8	73.3
Explanations	62.7	79.5	68.3	80.1	68.4	47.9	72.3
Examples	61.3	76.9	70.6	72.9	66.0	52.3	73.6
Roleplay	64.7	79.7	68.3	80.6	<i>65.5</i>	<i>47.1</i>	<i>70.7</i>
CoT	65.1	79.0	63.3	72.9	67.7	47.9	73.9
APO	67.9	80.7	70.0	75.0	69.2	55.4	74.4

Table 1: Accuracy (in %) of the hand-crafted prompts and of the best prompt obtained using automatic prompt optimization (APO) on each of the datasets and tasks described in Subsection 3.1. The highest accuracy achieved on a task is in bold, the lowest accuracy is in italics.

Inter-prompts performance variability

The accuracy⁴ on each task of our five types of prompts and of the best prompt generated using automatic prompt generation are reported in Table 1. We observe that the choice of prompt has a significant impact on performance for most tasks: depending on the task, the worst prompt results in between 14% (LibCon) and 47% (TE-off) more errors than the best prompt. This confirms our claim that careful prompt selection is crucial.

We also note that there is no "miracle" hand-crafted prompt that consistently outperforms all the others. Some heuristics and specific wordings work well on certain tasks, and poorly on others. This is in line with observations made in the literature; as an example, adding “Let’s think step by step” at the end of a prompt is shown to outperform adding the semantically similar “Let’s work this out in a step by step way to be sure we have the right answer” for the task studied in [YWL⁺24] (see Table 1), while the converse is true in [ZMH⁺22] (Table 7).

Automatic prompt optimization outperforms hand-crafted prompts

We now turn to the specific performances of Automatic prompt optimization (APO). APO beats hand-crafted prompts at all tasks but two; among those two, it is second best for TE-sent (with a score nearly equal to that of the best hand-crafted prompt), and is roughly at the level of the median of the hand-crafted prompts for TE-off⁵. This shows that automatically optimizing prompts results in consistently good performance without the need for prior testing of various hand-crafted prompts on a subset of the dataset of interest.

Interestingly, and in line with the remarks of the previous paragraph, the final prompts generated using APO do not differ much from the ones used to start the optimization process, despite the often large gap between their associated accuracies⁶. For TE-hate, for example, the initial prompt is “Classify the following message as hateful if it contains hate speech or non-hateful if it does not contain hate speech. Output only “hateful” or “non-hateful” without quotes” and the automatically optimized prompt is “Check for hate speech in the following message to determine if it is hateful, then classify it as either "hateful" or "non-hateful"”.

⁴We also recorded F1 scores, but do not report them, as they are very correlated with the accuracy and hence somewhat redundant.

⁵Note that TE-off has the smallest test set of all the tasks, and as a result the measured accuracies suffer from a greater variance.

⁶We report the best prompt returned by the optimization process for each task in the Appendix.

Is ChatGPT cheating?

Train set vs test set

	Simple	Explanations	Examples	Roleplay	CoT	APO
TE-hate, train	76.3	77.6	78.1	77.7	78.8	76.7
TE-hate, test	57.0	62.7	61.3	64.7	65.1	67.9
TE-emotion, train	76.9	75.9	74.5	77.2	76.3	81.0
TE-emotion, test	78.7	79.5	76.9	79.7	79.0	80.7

Table 2: Accuracy (in %) of the hand-crafted prompts and of the best prompt obtained using automatic prompt optimization (APO) on the train set and on the test set of TE-hate and TE-emotion datasets.

We have observed a curious phenomenon: as shown in table 2, all prompts result in much higher accuracy when tested on the training set of the TE-hate dataset than on its test set (where “training” and “test” refer to the split made by the creators of the dataset [BBF⁺19]). This is a priori surprising, as both subsets should be drawn from the same distribution, and as such our prompts should not perform any better on the training set than on the test set⁷. Our explanation is that the training set of TE-hate is part of the enormous amount of data on which GPT-3.5 Turbo, the model which we used, has been trained; as a result, it performs better on it than on the test set (which was either not included, or included in a different fashion). As the exact data on which the model was trained was not made public, we cannot confirm this hypothesis with absolute certainty. As a result, we can say that ChatGPT is “cheating” when labelling the training set of TE-hate, in the sense that its performance is superior to what it would be on a similar dataset to which it did not have access during its training (e.g. data that was produced after its training). This explains why we used the test splits of our datasets whenever possible, in the hope that as (we assume it to be the case) for TE-hate, they have not been included in the training set of the model, making for a fairer assessment of the capabilities of the method.

This problem makes it hard to predict future performances based on past experiments, as we have no guarantee as to what was and was not used in the models’ training. Note for example that the same gap in performance between the training set and the test set is not observable for TE-emotion. Ideally, one should only test the models on data that was produced after their training. This could however prove difficult and overly restrictive: curating useful datasets is a challenging and time-consuming task in itself, and most datasets available online are at least a few years old, thereby precluding all last generation LLMs.

As a side note, this phenomenon could explain why we have observed slightly lower accuracies on average than those reported in recent literature [HCvH24, GAK23, Tör23] on similar tasks: we suspect that the authors of these articles might not have been aware of this pitfall, and have taken no measures to circumvent it. It could also be due to intrinsic differences in difficulty between the datasets used, as we tried to include both easy and hard datasets for diversity (in particular, it has been noted before, e.g. in [HCvH24], that performance drops for very long texts, such as the articles from the AS-pol dataset). We cannot verify this hypothesis, as (most of) the datasets used in the references that we cite are not available anymore due to a policy change from Twitter.

5 Conclusion

This paper meant to illustrate the importance of prompt selection and the effectiveness of automatic prompt optimization within the context of text annotation tasks frequently used in the social sciences. In particular, our results suggest that the prompt optimization procedure described in Subsection 3.3 can be applied off-the-shelf to various tasks to achieve accuracies that equal or exceed those obtained using hand-crafted prompts. This prompt optimization method is made particularly easy by our user-friendly implementation⁸, which is freely available at <https://prompt-ultra.github.io/>.

⁷Note that there is no issue of overfitting here, as the prompts were not “trained” in any way on what we call the training set in this context.

⁸In particular, it allows for the replication of all of our experiments without the need for any coding.

Nonetheless, we have also seen that any automatic annotation procedure relying on last generation LLMs raises important questions of replicability due to their training processes (not to mention environmental and privacy-related concerns).

Among other potential future research directions, it would be particularly interesting to test whether LLMs tasked with labelling a corpus can give robust justifications for their choices – e.g. explain that they classified a given tweet as left-wing due to the presence of such and such opinions that are typically associated with the Left. Though more sophisticated solutions are conceivable (e.g. using the LLMs’ attention mechanisms, see [VSP⁺17]), directly asking the chatbot "Can you justify your decision?" would already make for an interesting experiment. What would also be extremely valuable is to be able to ask the LLMs to associate a (reliable) confidence score to its labels, so that human annotators can review the small percentage of labels the model was unsure of, and increase accuracy.

References

- [AAS⁺20] Mohd Zeeshan Ansari, M.B. Aziz, M.O. Siddiqui, H. Mehra, and K.P. Singh. Analysis of political sentiment orientations on twitter. *Procedia Computer Science*, 167:1821–1828, 2020. International Conference on Computational Intelligence and Data Science.
- [BBF⁺19] Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, Minnesota, USA, 2019. Association for Computational Linguistics.
- [BCCEAN20] Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa-Anke, and Leonardo Neves. TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification. In *Proceedings of Findings of EMNLP*, 2020.
- [BDSMGN20] Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. We can detect your bias: Predicting the political ideology of news articles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, EMNLP ’20, pages 4982–4991, 2020.
- [BEACC22] Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 258–266, Marseille, France, June 2022. European Language Resources Association.
- [BG24] Rick Battle and Teja Gollapudi. The unreasonable effectiveness of eccentric automatic prompts. *ArXiv*, abs/2402.10949, 2024.
- [BLHJ16] L.F. Barrett, M. Lewis, and J.M. Haviland-Jones. *Handbook of Emotions, Fourth Edition*. Psychology (The Guilford Press). Guilford Publications, 2016.
- [BMR⁺20] 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. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [DMB23] Nicolas Devatine, Philippe Muller, and Chloé Braud. An Integrated Approach for Political Bias Prediction and Explanation Based on Discursive Structure. In *Findings of the Association for Computational Linguistics (EACL 2023)*, pages 11196–11211, Toronto, Canada, July 2023. ACL: Association for Computational Linguistics.
- [GAK23] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. Chatgpt outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30):e2305016120, 2023.

- [GL12] JOSHUA S. GANS and ANDREW LEIGH. How partisan is the press? multiple measures of media slant*. *Economic Record*, 88(280):127–147, 2012.
- [HCvH24] Michael Heseltine and Bernhard Clemm von Hohenberg. Large language models as a substitute for human experts in annotating political text. *Research & Politics*, 11(1):20531680241236239, 2024.
- [JD16] Anuja P Jain and Padma Dandannavar. Application of machine learning techniques to sentiment analysis. In *2016 2nd International Conference on Applied and Theoretical Computing and Communication Technology (iCATccT)*, pages 628–632, 2016.
- [KGR⁺22] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.
- [LDC⁺22] Andrew Kyle Lampinen, Ishita Dasgupta, Stephanie C. Y. Chan, Kory Matthewson, Michael Henry Tessler, Antonia Creswell, James L. McClelland, Jane X. Wang, and Felix Hill. Can language models learn from explanations in context? *ArXiv*, abs/2204.02329, 2022.
- [LHM⁺23] Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, Zihao Wu, Lin Zhao, Dajiang Zhu, Xiang Li, Ning Qiang, Dingang Shen, Tianming Liu, and Bao Ge. Summary of chatgpt-related research and perspective towards the future of large language models. *Meta-Radiology*, 1(2):100017, 2023.
- [lib] Liberals vs Conservatives on Reddit. <https://www.kaggle.com/datasets/neelgajare/liberals-vs-conservatives-on-reddit-13000-posts>. Accessed: 2024-06-16.
- [LJW⁺22] Ruibo Liu, Chenyan Jia, Jason Wei, Guangxuan Xu, and Soroush Vosoughi. Quantifying and alleviating political bias in language models. *Artificial Intelligence*, 304:103654, 2022.
- [MBMSK18] Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. Semeval-2018 task 1: Affect in tweets. In *Proceedings of the 12th international workshop on semantic evaluation*, pages 1–17, 2018.
- [MPNR24] Fabio Motoki, Valdemar Pinho Neto, and Victor Rodrigues. More human than human: Measuring chatgpt political bias. *Public Choice*, 198(1):3–23, 2024.
- [NR13] M S Neethu and R Rajasree. Sentiment analysis in twitter using machine learning techniques. In *2013 Fourth International Conference on Computing, Communications and Networking Technologies (ICCCNT)*, pages 1–5, 2013.
- [Ope23] OpenAI. Gpt-4 technical report. 2023.
- [PP11] Marco Pennacchiotti and Ana-Maria Popescu. A machine learning approach to twitter user classification. In *Proceedings of the international AAAI conference on web and social media*, volume 5, pages 281–288, 2011.
- [pro] Prompt engineering guide. <https://www.promptingguide.ai/techniques/>. Accessed: 01-06-2024.
- [RBOP24] Stig Hebbelstrup Rye Rasmussen, Alexander Bor, Mathias Osmundsen, and Michael Bang Petersen. ‘Super-unsupervised’ classification for labelling text: online political hostility as an illustration. *British Journal of Political Science*, 54(1):179–200, 2024.
- [RFN17] Sara Rosenthal, Noura Farra, and Preslav Nakov. Semeval-2017 task 4: Sentiment analysis in twitter. In *Proceedings of the 11th international workshop on semantic evaluation (SemEval-2017)*, pages 502–518, 2017.
- [SDZ23] Kashun Shum, Shizhe Diao, and Tong Zhang. Automatic prompt augmentation and selection with chain-of-thought from labeled data. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12113–12139, Singapore, December 2023. Association for Computational Linguistics.

- [SM10] Carlo Strapparava and Rada Mihalcea. *Annotating and Identifying Emotions in Text*, pages 21–38. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [Tör23] Petter Törnberg. Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588*, 2023.
- [VSP⁺17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. 2017.
- [WR23] Maximilian Weber and Merle Reichardt. Evaluation is all you need. prompting generative large language models for annotation tasks in the social sciences. a primer using open models, 12 2023.
- [WWS⁺22] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Huai hsin Chi, F. Xia, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *ArXiv*, abs/2201.11903, 2022.
- [YCX⁺23] Junjie Ye, Xuanting Chen, Nuo Xu, Can Zu, Zekai Shao, Shichun Liu, Yuhan Cui, Zeyang Zhou, Chao Gong, Yang Shen, Jie Zhou, Siming Chen, Tao Gui, Qi Zhang, and Xuanjing Huang. A comprehensive capability analysis of gpt-3 and gpt-3.5 series models. *ArXiv*, abs/2303.10420, 2023.
- [YWL⁺24] Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers, 2024.
- [ZLS⁺24] Travis Zack, Eric Lehman, Mirac Suzgun, Jorge A Rodriguez, Leo Anthony Celi, Judy Gichoya, Dan Jurafsky, Peter Szolovits, David W Bates, Raja-Elie E Abdunour, et al. Assessing the potential of gpt-4 to perpetuate racial and gender biases in health care: a model evaluation study. *The Lancet Digital Health*, 6(1):e12–e22, 2024.
- [ZMH⁺22] Yongchao Zhou, Andrei Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers, 11 2022.
- [ZMN⁺19] Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 75–86, 2019.

A Brief overview of our browser-based service

We summarily describe our automatic text labelling service <https://prompt-ultra.github.io/>.

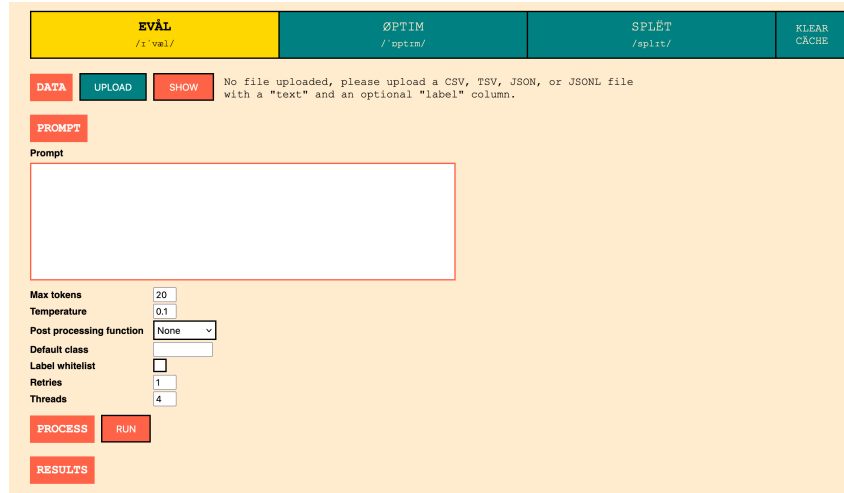


Figure 1: The EVÅL tab.

In the EVÅL tab (see Figure 1), you upload a dataset, which can be labelled or unlabelled, using the topmost button. You then enter a prompt and press the *Run* button. At that point, a pop-up window requires you to input your ChatGPT access key. Once you have done so, the prompt is applied to each entry of the dataset, and the resulting labelled dataset is output. If the original dataset was labelled, its labels and the predicted labels are compared and the resulting accuracy is computed.

Prompt optimization is performed in the ØPTIM tab (see Figure 2). You first upload a labelled dataset using the topmost button, and input a starting prompt. A macroprompt is then applied to generate successive generations of prompts, with the best prompts of each generation surviving to the next generation. A graph illustrates which prompt descends from which prompt of the previous generation. After a specified number of generations, the prompt whose performance was the best on the labelled dataset is output, as well as its accuracy score.

The SPLËT tab (see Figure 3) simply provides a convenient way to split a dataset in two.

Finally, the KLEAR CÄCHE button allows you to clear the LLM’s cache, and D.SKÖNNECT to erase the ChatGPT access key that you had entered.

B Optimized prompts

For each task, we report the best prompt generated by the prompt optimization process:

TE-hate “Check for hate speech in the following message to determine if it is hateful, then classify it as either "hateful" or "non-hateful".”

TE-emotion “Identify the emotion displayed in the following message as joy, anger, sadness or optimism.”

TE-sent “Label the emotion in the given message as positive, negative, or neutral.”

TE-off “Determine if the following message is offensive or non-offensive and provide the corresponding label.”

TML-sent “Categorize the sentiment in the following message as positive, negative, or neutral.”

AS-pol “Identify if the text below belongs to the left, center, or right categories.”

LibCon “Identify the text as either liberal or conservative.”

EVÁL /i-val/
ØPTIM /optim/
SPLËT /split/
KLEAR CÄCHE
D. SKÖNNËCT

DATA UPLOAD SHOW 324 records read from validation.jsonl

PROMPT

Prompt
 Write 0 if the text is negative, 1 if it's neutral and 2 if it's positive.

PARAMETERS

Max tokens:
 Temperature:
 Post processing function:
 Default class:
 Label whitelist:
 Retries:
 Threads:

PARAMETERS

Metaprompt
 Generate a variation of the following instruction while keeping the semantic meaning and the output format:
 {prompt}
 Output only the new instruction.

Max prompt tokens:
 Population size:
 Generation multiplier:
 Generations:
 Temperature:

OPTIMIZE RUN STOP

118/324

```

graph LR
    A((0.605)) --- B((0.590))
    A --- C((0.627))
    C --- D((0.627))
    C --- E((0.630))
    C --- F((0.633))
    C --- G((0.657))
    
```

RESULTS

Best prompt:
 Please assign the value of 0 to negative text, 1 to neutral text, and 2 to positive text.
Score: 0.6574074074074074

Figure 2: The ØPTIM tab.

EVÁL /i-val/
ØPTIM /optim/
SPLËT /split/
KLEAR CÄCHE
D. SKÖNNËCT

DATA UPLOAD SHOW 324 records read from my_dataset.jsonl

PARAMETERS

Seed:
 Number of samples:

SPLIT RUN

RESULTS

DOWNLOAD SHOW my_dataset-split1.jsonl: 100 records
DOWNLOAD SHOW my_dataset-split2.jsonl: 224 records

Figure 3: The SPLËT tab.