Studies in Communication Sciences 24.1 (2024), pp. 161–169

Botlitica: A generative AI-based tool to assist journalists in navigating political propaganda campaigns

Elena Musi*, University of Liverpool, Department of Communication and Media, UK
Edgar Everardo Garcia Aguilar, AXA UK, Data Science Team
Lorenzo Federico, Luiss University, Datalab, Italy
*Corresponding author: elena.musi@liverpool.ac.uk

Abstract
The hype on generative AI has raised concerns about the spread of disinformation but also opened up new opportunities for hybrid journalism. The proliferation of political propaganda campaigns spread across digital media during election periods constitutes a challenge for journalists who struggle to exercise information gatekeeping. AI-based tools can in principle help, but journalists are resistant to using them since they are sceptical about their compliance with news values and their scarce user-friendliness. To cope with such issues, we present *Botlitica,* a GPT3-based chatbot able to answer users' questions according to the information shared on different social media by a given political party which is, thus, embodied by the AI agent in the conversation. The back-end and front-design of the chatbot are devised to privilege *transparency.* We report the results of a preliminary evaluation of *Botlitica* which show that the tool fastens journalists' capabilities to navigate propaganda campaigns inducing them to exercise critical thinking.

Keywords
disinformation, generative AI, political communication, human computer interaction, fact-checking

1 Introduction

Keeping up with the proliferation of information across digital media constitutes one of the major challenges in the current information ecosystem. In the context of snap political elections, such as the Italian ones which happened on the 25th September 2022, that of navigating the Networked Society is a necessary endeavour i) for political parties to spread propaganda and ii) for citizens to get informed and shape their decision-making processes accordingly. At the same time, to provide accurate information about political propaganda campaigns, let alone to evaluate them, journalists and communication strategists are called to keep up with a proliferation of messages sparse across social media platforms and hindered by misinformation. To cope which such a scenario, *data journalism* – news making based on the automatic analysis of large datasets – has become one of the five most important journalistic innovations developed in Austria, Germany, Spain, Switzerland, and the United Kingdom from 2010 to 2020 (Meier et al., 2020). Computational social science techniques that allow the automatic retrieval of information (e. g., classification systems of propaganda rhetorical strategies) have proven to be useful analytic tools (Liew & Mueller, 2022; Petridis et al., 2023). However, such tools are rarely used by practitioners beyond academia since they raise doubts as to their transparency and they are poorly user-friendly.

Leveraging on advances in generative AI, we pilot a GPT-3 (Brown et al., 2020) based chatbot in Italian *(Botlitica)* able to answer users' questions according to the information shared on different social media by a given political party which is, thus, embodied by the AI agent in the conversation. The interface shows the set of social media messages used by the system as a source to come up with the answer. Besides guaranteeing *transparency*, such a design allows communication practitioners to grasp information across topics and platforms in a conversational format, while retrieving relevant data which can be further analysed. To test how the tool

https://doi.org/10.24434/j.scoms.2024.01.4270

would be perceived by communication practitioners, we trial its use in a focus group with 18 journalists. The journalists had to simulate being editorialists having to write a piece commenting upon the parties' propaganda campaigns. Particular attention is devoted to the usefulness of the tool for the identification of rhetorical strategies (e.g., lack of consistency in explaining political agendas; vagueness of expression etc.) in propaganda campaigns. Previous research has, in fact, shown that political discourse is hindered by fallacious arguments which are not necessarily intentional (Musi, Aloumpi, Carmi, Yates, & O'Halloran, 2022) but harmful since highly misleading for the wider public.

2 *Botlitica* design

This section is devoted to the presentation of Botlitica infrastructure from a software and a user perspective.

2.1 Back-end design

GPT-3 (Brown et al., 2020) is a pre-trained autoregressive language model with 175 billion parameters that can generate human-like text. Since it has been trained on an encompassing dataset (Common Crawl, WebText2, Books1/Books2 and Wikipedia) which well approximates language use on the Internet, it performs very well on a variety of language tasks: it can, for instance, generate quality poetry and prose, even in the style of an existing author. However, if the author revolutionised their style and content starting from 2022, GPT-3 would not be able to mimic such a turn, lacking the knowledge (training data). Similarly, if a political party changed views over a matter from one election round (2018 for Italy) to another (2022), it would not be captured by GPT-3.

As a first step to make GPT-3 up to date with the rapidly changing Italian political scenario (including new political alliances), it was, thus, necessary to fine tune the model using the OpenAI *CLI* commands, after having obtained an API key from OpenAI. As a pilot, we decided to focus on three major parties (*Fratelli d'Italia*, *Partito Democratico* and *Terzo Polo*) and Twitter and Facebook official accounts as social media sources. Besides reaching different audiences, the two social media platforms are the main ones that feature all three parties active.

We collected through the *Twitter API* and *Crowdtangle* all the tweets / posts shared through the main official Twitter handles/FB posts associated with the parties in the period from 21st July 2022 (government fall) to 15th September 2022. We then pre-processed the dataset i) discarding all the tweets and posts with less than 10 words and made up of a single URL, ii) deleting the hash symbol from hashtags while integrating them as part

Table 1 Sources composing the training datasets for fine tuning

	Fratelli d'Italia	Partito Democratico	Terzo Polo
Twitter	@fratelliditalia	@pdnetwork	@Azione_it
	@giorgiameloni	@nzingaretti	@CarloCalenda
	@Ignazio_LaRussa	@EnricoLetta	@MatteoRichetti
	@guidocrosetto	@AndreaOrlandosp	@PastorellaGiu
	@DSantanche	@serracchiani	@Enrico__Costa
	@Giov_nazionale	@simonabonafe	@Filippo_Rossi
	@LucioMalan	@pierofassino	@msgelmini
	@Alberto_cirio	@sbonaccini	@mara_carfagna
	@marcomarsilio	@pbersani	@ItaliaViva
	@FidanzaCarlo	@nomfup	@MatteoRenzi
	@fabiorampelli	@graziano_delrio	@marattin
	@Donzelli	@micheleemiliano	@lucianonobili
	@raffaelefitto	@mariannamadia	@meb
	@francescolollo1	@robertapinotti	@ivanscalfarotto
	@isabellarauti	@PaoloGentiloni	@TeresaBellanova
	@adolfo_urso	@Pierferdinando	@bobogiac
Facebook	Noi_con_Giorgia_Meloni	Partito Democratico	Carlo_Calenda
	Giorgia_Meloni_Presidente	Enrico Letta	Azione
	Giorgia_Meloni	Deputati PD	Matteo_Renzi

Figure 1: Botlitica onboarding page

Source: https://github.com/musielena.

of a sentence where possible (e.g., “#VotaFdi #ElezioniPolitiche2022” → “Vota FdI alle Elezioni Politiche 2022”) and otherwise remove them, iii) removing the reporting clause in case of direct speech. Finally, we manually picked a sample of 300 tweets and 300 Facebook posts per each party to guarantee a variety of topics and clarity while manually checking the relevance of the message.

We then assembled tweets and Facebook posts in a single training dataset per party in *JSON* format. We fine-tuned a *davinci* model for each of the parties using embeddings (OpenAI Embeddings API), to create a question-and-answer bot which is based on additional knowledge. We selected *davinci* since it is more capable than the default model (*curie*). We ruled out a simple fine-tuning model (with tweets and Facebook posts as additional training data), since it was giving in output hallucinations, e.g., answers that were not at all in line with the party view or not relevant to the questions. This is not surprising since the process of fine-tuning adjusts the weight of the base model, which however works well “when the answer is contained within the paragraph, however if the answer isn’t contained, the base models tend to try their best to answer anyway, often leading to confabulated answers” (OpenAI documentation: http://tinyurl.com/24str9m5). On the other side, embeddings constrain the answers to a set of relevant contexts (the social media data in our case): An embedding is created for each of the social media posts; the user’s query section of the social media data that are similar to the prompt are then selected by the system; a dynamic prompt with this information is created to answer the question. It took a few iterations (fine-tune a model with embeddings and test it on the *Openai playground*) to come up with a well-performing model. As a test prompt we used the following: “Quali sono i punti principali del programma politico di Fratelli d’Italia?” (“What are the main points of the political program of the Brothers of Italy?”).

2.2 Front-end design

The next step was to customise the chatbot using the fine-tuned model in which the AI assistant embodies a given political party’s view expressed on social media. Drawing from the Q&A preset from OpenAI, we input as a prompt a short bot’s identity description coupled with an example of answer / questions for each politician:

> Example: “Rispondi alla domanda come se fossi [the candidate].\n\n”+question+‘\\n’ (“Answer to the question as if you were [the candidate].”)

Figure 2: Example of user interaction with *Botlitica*

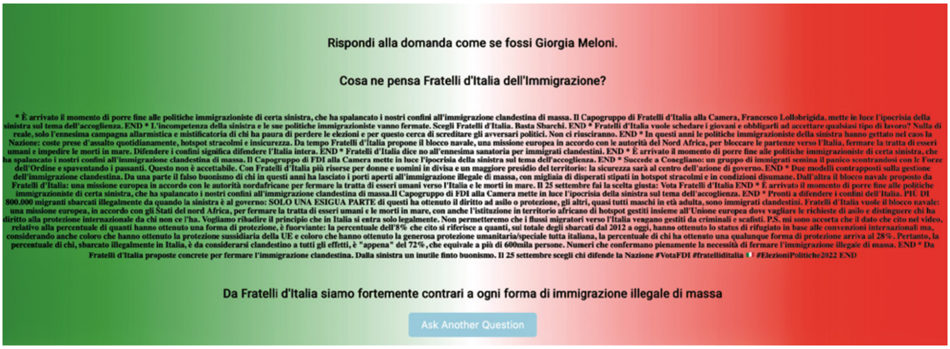

Source: https://github.com/musielena.

> Human: Buongiorno. Perché dovrei votare per Fratelli d'Italia? ("Goodmorning. Why shall I vote for Brothers of Italy?")
> AI: Perché FdI è l'unico partito pronto a combattere la povertà, investendo sul lavoro ("Because FdI is the only party ready to fight poverty by investing in work")
> Human: Cosa ne pensi dell'immigrazione? ("What do you think about immigration?")

After having checked the soundness of the completions in the playground, we then used OpenAI's API to generate GPT-3 responses in Python. This allowed us to implement a chat. log function to ensure that the bot would remember the dialogue history.

To allow users to easily interact with the chatbot, we made the bots respond in real time using the Twilio SMS API. As a webhook to route the messages we built a Flask app and configured Twilio accordingly. As a web service to deploy the Flask App we opted for the AWS (Amazon Web Services) console. The user interface contains a short description of *Botlitica* as well as a three-tiered candidate choice among the representatives of the three parties *Fratelli d'Italia*, *Partito Democratico* and *Terzo Polo*.

After having chosen one of the candidates, the user is free to ask any question, getting in output a conversational answer coupled with the set of tweets / Facebook posts used as a knowledge base to build the answer.

In Figure 2 the journalist is asking Giorgia Meloni what the party Brothers of Italy thinks about immigration. The answer ("Brothers of Italy strongly opposes any form of mass illegal immigration") is drawn from the set of tweets and Facebook posts reported above it. The printed social media posts do not represent all the messages chosen as context by the embeddings, but the closest to the question based on cosine similarity due to a restriction in the number of tokens that the model can take as a prompt (2049 tokens). In discourse terms, the set of social media posts allows the journalists to access the sources of information used by GPT-3 to come up with the answer. This builds common ground while allowing the journalist to retrieve relevant information in a fast-paced manner, while exercising epistemic vigilance towards the synthetic output.

3 *Botlitica* evaluation

This section reports on the results of a focus group with practitioners to evaluate the tool.

3.1 Focus group set up

To investigate users' perceptions about *Botlitica*, we organised a 2 hours focus group with 18 journalists attending the School of Journalism at Luiss Guido Carli University. Before the focus group, the journalists attended a 1.30 hours lecture about Rhetorical Discourse for propaganda purposes. First, they were introduced to the Pragma-dialectic notion of strategic manoeuvring (Van Eemeren & Houtlosser, 2000), intended as a set of discourse strategies to achieve at once dialectical reasonableness and rhetorical effectiveness. The three main forms of strategic manoeuvring were defined as follows and explained with the aid of examples:

- *Topical choice* refers to the themes that arguments rely upon (e.g., choosing an event rather than another one to make an argument)
- *Audience adaptation* involves framing one's moves in a perspective that agrees with the audience (e.g., referring to personal experiences)
- *Presentational devices* concern the selection of stylistic features which persuade the audience (tone of language, use of emotional terms to create empathy)

The second half of the lecture was devoted to the presentation of the "Propaganda Analysis project" (Da San Martino, Barron-Cedeno & Nakov, 2019), which identified a core set of propaganda techniques (e.g., *Red Herring*, *Strawman*, *Black and White* fallacy), drawing from the annotation of 451 news articles from 48 news outlets, both propagandistic and non-propagandistic according to MediaBias. Each of the 20 propaganda techniques was briefly discussed and the journalists were provided with the link to a cheat sheet containing definitions and examples.[1]

During the first half hour of the focus group, the participants were provided with a brief explanation of the functionalities of the tool, the basic notions underlying generative AI and how the data emerging from the interaction with the tool would be used. The participants were then asked to simulate being a journalist having to write an opinion piece during the snap elections electoral campaign. Each of them was provided with a link to a folder containing one of the party's social media data (same datasets used to fine-tune *Botlitica*) and with access to *Botlitica*. The task was the following:

- Step 1: Analyse the social media datasets to investigate the propaganda strategies of the assigned party (30 Min)
- Step 2: Complete *Qualtrics* questionnaire 1 (15 min)
- Step 3: Play with *Botlitica* asking questions to the assigned party (different from the one in step 1) (30 Min)
- Step 4: Complete *Qualtrics* questionnaire 2 (15 min)

In such a way we obtained a *test group* of 18 answers (6 per political party) after having used *Botlitica* (from now on TG) and a control group of 18 answers gathered from the standard data analysis of published official propaganda messages (from now on CG). For comparison purposes, the 5 questions in Questionnaire 1 and 2 were the same:

Q1: What party are you evaluating?
Q2: Overall, how persuasive is the campaign?
Q3: Why do you think it is persuasive or not?
Q4: What are the propaganda strategies that you encountered?
Q5: Provide examples of the encountered propaganda strategies that you remember.

3.2 Results

Since one of the participants did not save the answers to the questionnaire, and another one just answered "0" to all slider questions in Q2 and Q3, the sample was reduced to 16 participants for both the test and the control group with 6 sets of answers for *Terzo Polo* and 5 answers each for *Fratelli d'Italia* and *Partito Democratico*.

As to perceived persuasiveness, the mean values show that the propaganda campaigns were considered more persuasive when accessed through *Botlitica*: The mean value for the TG amounts to 51.4, while the mean value for the CG to 40.8, while the standard deviations are similar (22.6 vs 26.2), suggesting that the trend is shared across participants. This is not surprising: Conversational interfaces resulted to be more persuasive than mediated ones (e.g., text-based) in experiments held across domains since the human-likeness induced by the conversational format increases the trust level (Schulman & Bickmore, 2009). Furthermore, it is reasonable to assume that the turn-taking format nudges the journalist to keep up with the conversation rather than

1 Cheatsheet available at: https://propaganda.qcri.org/annotations/definitions.html.

Table 2 Perceived persuasiveness across the forms of strategic manoeuvring in the test group (TG)

Form	Mean	Std Deviation	Variance	Count	p-value
Topical potential	46.7	25.3	640.1	16	0.879
Audience demand	49.9	22.2	497.0	16	0.691
Presentational devices	47.0	25.0	625.6	16	0.070

Table 3 Perceived persuasiveness across the forms of strategic manoeuvring in the control group (CG)

Form	Mean	Std Deviation	Variance	Count
Topical potential	46.6	25.0	626.1	16
Audience demand	47.3	23.4	549.4	16
Presentational devices	34.6	24.4	595.7	16

Table 4 Propaganda strategies recognized by the test group (TG) and the control group (CG)

Propaganda strategies	TG %	TG Count	CG %	CG Count
Appeal to authority	5.5	5	2.7	2
Appeal to fear/prejudice	5.5	5	8.2	6
Bandwagon	1.1	1	2.7	2
Black-and-white Fallacy, Dictatorship	4.4	4	6.9	5
Causal Oversimplification	11.0	10	11.0	8
Doubt	3.3	3	0.0	0
Exaggeration or Minimisation	6.6	6	9.6	7
Flag-waving	5.5	5	5.5	4
Glittering Generalities (Virtue)	0.0	0	4.1	3
Loaded Language	2.2	2	4.1	3
Misrepresentation of Someone's Position (Straw Man)	2.2	2	5.5	4
Name calling or labeling	2.2	2	4.1	3
Obfuscation, Intentional vagueness, Confusion	6.6	6	1.4	1
Presenting Irrelevant Data (Red Herring)	4.4	4	4.1	3
Reductio ad hitlerum	2.2	2	1.4	1
Repetition	8.8	8	5.5	4
Slogans	11.0	10	8.2	6
Smears	2.2	2	0.0	0
Thought-terminating cliché	2.2	2	1.4	1
Whataboutism	13.2	12	13.7	10
Total	100	91	100	73

taking time to reflect upon fallacious rhetorical aspects of a message. These results suggest that propaganda strategies communicated through chatbots might be more effective in persuading the public, calling for the need for critical thinking of the role played by generative-AI fuelled agents implemented in the future. Zooming into the 3 forms of strategic manoeuvring the rates of perceived persuasiveness are presented in Table 2 and 3.

Both in the TG and the CG the audience demand is perceived as the most persuasive features; presentational devices are, instead, considered by far the least persuasive by the CG while this is not the case for the TG. This confirms the engaging role played by language in a conversational, rather than written format which creates a feeling of reciprocation and mutual understanding not achievable through tweets and Facebook posts. The latter might be tailored towards an audience with respect to whom the speaker does not feel included.

Turning to question 4, the most apparent result is that the TG recognised a larger number of rhetorical strategies (91) compared to the CG (73). This result was expected since the main goal of *Botlitica* is that of streamlining the process of information retrieval customised by the journalist more flexibly, compared to standard query tools based, for example, on keywords' search. Looking at the breakdown per single propaganda strategy (see Table 4), *repetitions* and *vagueness* are the strategies where access to *Botlitica* seem to especially fastened the identification process: Allowing for a cumulative search for relevant contexts (tweets) to answer a question, *Botlitica* allows the journalists to recognise the *dispositio* and the *elocutio* patterns of social media, otherwise scattered across days and platforms, and thus concealed. On the other hand, the trend is reversed when it comes to "language strategies": The CG has identified more cases of *Exaggeration or Minimisation, Glittering Generalities, Loaded Language and Name Calling* with comparison to the TG. This might be due to the catering role played by the AI-answer which, when showing more neutral tones compared to the contexts, might defocus the attention to the rhetorical role played by lexical aspects.

As to the other strategies, which require an inferential process to be interpreted, the differences in numbers between TG and CG are tiny and anyways not significant in such a small sample. The same limitation applies to the analysis of Q5 which received 5 answers by the CG and 6 by the TG. The identified strategies in both cases were correct a part from 1 stemming from the CG („Choose between *us* and *them*") where an *appeal to fear* was identified instead of a *Black and White* fallacy. Regardless of the group, the *Black and White Fallacy* and *Strawman* were the most remembered strategies.

4 Conclusion

Propaganda campaigns under election periods are challenging to navigate for journalists due to the proliferation of information sparse across digital media. Journalists' gatekeeping function is, however, paramount to avoid the phenomena of political polarisation and disinformation / misinformation which are fuelled by digital media affordances. Computational tools leveraging Generative Artificial Intelligence such as GPT-3 can help facilitate both information retrieval and analysis to assist the news-making process. However, there seems to be a gap between the number of available GPT-3 applications and their actual use in the newsrooms mainly due to i) difficulties in learning how to use these tools, and ii) worries about their adherence to news values. As to the latter, the black-box nature of the processes underlying their outputs constitutes a major concern raising, for example, questions about accountability. The hallucinations (untruthful outputs) which feature the outputs of Large Language Models (Ouyang et al., 2022) crucially undermine their trustworthiness. While some of these hallucinations depend on the time-limitations of large pre-trained models (up-to-date till 2021 in the case of GPT-3), others cannot be predicted and are difficult to explain, calling for a degree of transparency still not available. To address these issues, we developed *Botlitica*, a GPT-3 based chatbot in Italian which allows to answer users' questions according to the information shared by a given political party which is, thus, embodied by the AI agent in the conversation. We pilot *Botlitica* with social media data (Twitter and Facebook) published by the official accounts of three major parties (*Fratelli d'Italia*, *Partito Democratico* and *Terzo Pol*o) during Italian snap elections.

The design of *Botlitica* is novel both at back-end and front-end level. At the back-end level the model is fine-tuned through embeddings rather than standard fine: As a result the social media posts are interpreted by the model as a privileged knowledge base to build an answer, reducing hallucinations. At the front-end level, journalists are provided not only with a synthetic conversational answer to their questions, but also with the set social media messages used by the system as context to draw the answer. In such a way, journalists can retrieve information relevant to complex queries, while maintaining epistemic vigilance with respect to the synthetic outputs. We believe that such a design advances *transparency*, engaging the user in critical discussion with the AI over an issue (question asked): Even though the inferences that brought *Botlitica* to craft an answer remain covert, the sources of information are disclosed, creating common ground knowledge with the user (journalist). The journalist can then agree or disagree with the interpretation provided by *Botlitica* (and expressed in its answer) which might, in turn, offer hints about the misleading nature of propaganda messages. In other words, even though *Botlitica*'s ways of reasoning might be fallacious, the risk of evading the burden of proof which characterises Generative Artificial Intelligence is overturned.

To evaluate the usefulness of *Botlitica* for journalists, we carried out a focus group where 18 journalists had to navigate the digital propaganda campaigns with or without *Botlitica* and complete a questionnaire meant to assess their analytic skills with respect to the propaganda discourses as instances of strategic communication. The results show that *Botlitica* has significantly fastened journalists' skills to identify propaganda strategies across the board. Overall, tools such as *Botlitica* with a transparency-centred design, could help second generation Newsroom Innovation Labs (Cools, Van Gorp, & Opgenhaffen, 2022), evolving *vis*-à-*vis* journalistic values.

Conflict of interest

The authors declare no conflict of interest.

References

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, *33*, 1877–1901.

Chavan, T., & Kane, A. M. (2022, December). Large language models for multi-label propaganda detection. In *Proceedings of the The Seventh Arabic Natural Language Processing Workshop (WANLP)* (pp. 515–519). https://doi.org/10.18653/v1/2022.wanlp-1.60

Cools, H., Van Gorp, B., & Opgenhaffen, M. (2022). New organizations, different journalistic roles, and innovative projects: How second-generation newsroom innovation labs are changing the news ecosystem. *Journalism Practice*. Advance online publication. https://doi.org/10.1080/17512786.2022.2097940

Da San Martino, G., Barron-Cedeno, A., & Nakov, P. (2019, November). Findings of the NLP4IF-2019 shared task on fine-grained propaganda detection. In *Proceedings of the second workshop on natural language processing for internet freedom: Censorship, disinformation, and propaganda* (pp. 162–170). https://doi.org/10.18653/v1/D19-5024

Da San Martino, G., Yu, S., Barrón-Cedeno, A., Petrov, R., & Nakov, P. (2019, November). Fine-grained analysis of propaganda in news article. 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)* (pp. 5636–5646). https://doi.org/10.18653/v1/D19-1565

Liew, A., & Mueller, K. (2022). *Using large language models to generate engaging captions for data visualizations*. arXiv preprint. Retrieved from https://arxiv.org/abs/2212.14047

Martino, G. D. S., Cresci, S., Barrón-Cedeño, A., Yu, S., Di Pietro, R., & Nakov, P. (2020). *A survey on computational propaganda detection*. Retrieved from https://arxiv.org/abs/2007.08024

Meier, K., Schützeneder, J., García Avilés, J. A., Valero-Pastor, J. M., Kaltenbrunner, A., Lugschitz, R., ... Saner, M. (2022). Examining the most relevant journalism innovations: A comparative analysis of five European countries from 2010 to 2020. *Journalism and Media*, 3(4), 698–714. https://doi.org/10.3390/journalmedia3040046

Musi, E., Aloumpi, M., Carmi, E., Yates, S., & O'Halloran, K. (2022). Developing fake news immunity: Fallacies as misinformation triggers during the pandemic. *Online Journal of Communication and Media Technologies, 12*(3), 1–18. https://doi.org/10.30935/ojcmt/12083

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., ... Lowe, R. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems, 35*, 27730–2774. Retrieved from https://proceedings.neurips.cc/paper_files/paper/2022/file/

b1efde53be364a73914f58805a001731-Paper-Conference.pdf

Petridis, S., Diakopoulos, N., Crowston, K., Hansen, M., Henderson, K., Jastrzebski, S., … Chilton, L. B. (2023, April). AngleKindling: Supporting journalistic angle ideation with large language models. In *Proceedings of the 2023 CHI conference on human factors in computing systems* (pp. 1–16). https://doi.org/10.1145/3544548.3580907

Schulman, D., & Bickmore, T. (2009, April). Persuading users through counseling dialogue with a conversational agent. In *Proceedings of the 4th international conference on persuasive technology* (pp. 1–8). https://doi.org/10.1145/1541948.1541983

Van Eemeren, F. H., & Houtlosser, P. (2000). Rhetorical analysis within a pragma-dialectical framework. *Argumentation, 14*(3), 293–305.