

# Predicting Perceived Agency in News Headlines

Igor Ryazanov

Umeå University

Department of Computing Science

igorr@cs.umu.se

## Abstract

This paper describes work in progress on computation analysis and prediction of perceived agency attribution in news articles. Perceived agency and its manipulation through linguistic means is a key component of a number of highly salient narratives, such as the dehumanisation of minoritised groups and the anthropomorphism of technology. We propose predicting agency in the news headlines as a computational task. To this end, we will annotate a pair of topical datasets, and analyse the predictive power of RoBERTa-based contextual embeddings.

## 1 Introduction

News headlines serve the combined purpose of attracting attention and summarising the content of articles and are often a convenient subject for studies of media coverage in general. Computational tasks for news articles often focus either on the broader NLP tasks, such as stance or emotion detection or choose a specific narrow narrative to analyse (e.g. framing of gun violence in the US: Liu et al. (2019), Tourani et al. (2021)). This paper proposes to focus instead on a mechanism that plays an important role in a variety of different narratives by investigating whose agency headlines communicate. We consider *perceived agency*, which represents readers' interpretation of the distribution of agency between the entities involved in the situation the headline describes. We see this type of agency as a low-level contributor to a number of impactful news narratives: responsibility placement in crime coverage ('Collision between bike and car' vs. 'Cyclist slams into car door', Minnema et al. (2022a)), anthropomorphism in technological coverage ('A self-driving truck handles long haul freight...' vs. 'A driverless truck is programmed to transport long haul freight...', Inie et al. (2024)) and a number of others. The perceived agency of these entities does

not necessarily correspond to a stance taken by the author and does not necessarily place responsibility for the situation on any of the entities but facilitates some of these interpretations. As such, it can be seen as an important linguistic component for building narratives and possibly a proxy for the power relations the publication chooses to highlight.

In this study, we annotate and train our model on two datasets of news headlines. One covers a broad range of news coverage, while the other focuses on a narrow topic: reporting on AI-related technologies. Our research questions are then as follows:

- **Q1:** Do human annotators produce consistent annotations of perceived agency?
- **Q2:** Can a computational model trained on such annotations reliably predict perceived agency based on contextual embeddings?
- **Q3:** Does the performance transfer between domains?

If the experiments prove successful and the task of predicting perceived agency is feasible and computationally solvable, the approach should prove valuable for the computational analysis of framing and narratives, as well as for the practical language analysis problems arising in digital media studies.

## 2 Background

In the body of literature dedicated to computational analysis of news coverage and headlines in particular there are a number of studies focusing on narrative elements that can be connected to readers' perceptions of agency.

The closest is, perhaps, the computational detection of perspectivisation and responsibility framing (Minnema et al., 2022b). The notion of responsibility clearly overlaps with agency but differs in that it itself has to be further defined within the

appropriate context. For example, in their study of responsibility perception in femicide reporting, [Minnema et al. \(2022a\)](#) define responsibility for the annotator through the combination of three aspects: *focus*, *cause* and *blame*. *Focus* (‘does the sentence focus on the agent or on something else?’), arguably can be seen as loosely corresponding to the perceived agency, while *blame* and *cause* are contextual to the criminal and, potentially, moral responsibility. It stands to argue that the perceived agency component can be decoupled from the notion of responsibility to make such analysis more generalisable. From the readers’ perspective, perceived agency can be inferred from the specific phrasing, while establishing responsibility requires a separate step of reasoning of invoking a much broader context (e.g. what do I, the reader, consider to be a cause?).

Another relatively commonly studied type of narrative closely related to perceived agency is the anthropomorphising language. Anthropomorphising language describing technological systems has attracted more attention recently with the advances of generative large language models. For instance, [Abercrombie et al. \(2023\)](#) warn against the dangers of extensive anthropomorphism of dialogue systems. It has also been considered as a factor affecting consumers’ trust in technology ([Inie et al., 2024](#)). The findings of [Cheng et al. \(2024\)](#) also suggest anthropomorphism becoming more present in the scientific literature over time, though still at a lower level than in downstream news coverage. Anthropomorphism in news articles has been shown to come in many forms ([Ryazanov et al., 2024](#)), but many of them can be tied directly to the notion of agency. Of the four classes of anthropomorphising language [Inie et al. \(2024\)](#) identify, one is ‘describing the machine as an agent of an action’ while two others (portrayal of the machine as a cogniser and as a communicator) are also inherently related to assigning a degree of agency to the technology.

The opposite type of narrative—the dehumanisation of human beings—similarly often relies on the manipulation of perceived agency. For instance, discursive techniques to humanise or dehumanise migrants can draw on agency ([Kirkwood, 2017](#)): portraying incoming migrants as independent agents (e.g. asylum seekers) in opposition to more passive roles (refugees) naturally affects public perception and can influence the assumptions underlying political decisions ([Sajjad, 2018](#)). However, similarly to anthropomorphising language,

dehumanisation is not limited to agency manipulation. For instance, in their dataset of dehumanising language, [Engelmann et al. \(2024\)](#) identify six categories: negative evaluation of a target group, denial of agency, moral disgust, animal metaphors, and objectification. Of these, only denial of agency and objectification can be tied directly to agency, while others use different mechanisms.

In all three cases, the (manipulation of) perception of agency is one of the important mechanisms through which narratives and framings are communicated. However, to our knowledge, this mechanism has not been studied separately through computational methods, which is the central goal of this study.

### 3 Perceived Agency

We define perceived agency as the agency the readers project on other entities while interpreting a description of a situation. Importantly, we use the term *perceived* to highlight that the level of agency in question is not necessarily the same as implied by the author because the reader does not operate in the same context. This matters, in particular, when working with news headlines that are often scanned without opening the article. Some examples of this are shifting agency for a humorous effect that is not obvious without context or the use of technical terms, the precise definition of which is not expected to be known by the broader public. The latter is exemplified by increasing AI anthropomorphism from scientific articles to news publications ([Cheng et al., 2024](#)). From the annotation perspective, considering agency attribution on the readers’ side means that human annotators are told to annotate their interpretation of a headline without explicitly trying to infer the author’s intention.

For perceived agency to be successfully annotated and used as a target value for a computational model, assumptions that (a) humans are consistent in describing the relative agentivity of entities in a given situation, (b) these levels of agentivity are largely determined by word choice and sentence structure and do not require additional context. Verifying (a) is an important part of this study, while (b) should be satisfied in the domain of news media as news stories tend to be relatively self-sufficient; summarising the article is one of the main purposes of news headlines, along with attracting attention.

Table 1: The news datasets used to sample headlines.

Dataset	Sources	Time period	Unique articles
General coverage	Open-access British news media	May 2022 - Aug 2022	57,996
AI coverage	Global open-access English-language media (UK, US, India)	June 2022 - May 2023	6055

### 3.1 Datasets

At the current stage, we plan to annotate headlines sampled from two English-language news datasets (Table 1). The first dataset is a collection of articles from British media over the summer of 2022 and serves as the domain unspecific benchmark for annotation and model performance. The second dataset is focused on a single topic—AI—and includes a full year of coverage in the most visited English-language news media. Both are self-collected and further described in our previous work: [Ryazanov et al. \(2024\)](#) and [Ryazanov and Björklund \(2023\)](#) respectively.

### 3.2 Annotation

We propose the following annotating process, which is currently in progress (An example shown in Table 2):

1. Annotators are given an explanation of their task: ‘Given this headline describing a situation, who or what affects the situation and to what extent?’
2. The respondents are shown an explanation of the five-point scale of levels of agency: ‘Full agency’ (5), ‘High level of agency’ (4), ‘Medium level of agency’ (3), ‘Low level of agency’ (2), and ‘No influence on the situation’ (1). For each level, examples are provided, which is also our motivation for limiting the scale to five.
3. For each noun phrase in a headline, the respondents answer a multiple-choice question about the level of agency the entity has in the situation (‘How would you describe the agency that NP demonstrates in this situation?’).
4. Each headline is annotated by at least two annotators. To aggregate the annotations for each headline, they are converted from categories to numerical values in the  $[0;1]$  interval and scaled to add up to 1.

Importantly, the resulting values are relative to situations described in each headline and, therefore, cannot be compared directly.

### 3.3 Predicting Perceived Agency

After testing the annotators’ agreement to answer **Q1**, we will proceed to the text step. The initial setup for our experiments will be training a regression model based on contextual embeddings from RoBERTa ([Liu, 2019](#)) that would predict the agency level value given a sentence (headline) and a noun phrase (an entity that may display agency). We aim to train such models individually on headlines from both datasets to evaluate and compare the performance over a general coverage domain to a more narrow topic. The choice of a pre-trained RoBERTa as the baseline model is largely due to its popularity. Since showing viability is the primary goal of the experiments, even indirect comparisons with the performance of BERT-based models on related tasks (e.g. [Minnema et al. \(2022a\)](#)) are a valuable point of reference.

## 4 Discussion and Limitations

This work is very much in progress, and the next steps can still be adjusted. In particular, depending on how successful the annotations and initial model trials are, we will consider other topical headline datasets to broaden the comparison. Our ambition is to demonstrate that agency perception can contribute to building a wide variety of narratives, and if that is the case, a model capable of predicting agency levels can serve as a versatile tool in media analysis. A secondary goal is, therefore, determining which model type is most successful at this task, which means further experiments with other training and inference setups, e.g. with the generative large language models.

It is worth acknowledging that our approach has a number of limitations, largely due to the difficulty of annotation. For instance, our annotators consider the headlines without knowing the source website, which is not how interactions with news normally play out. Our definition of agency is also relatively limited and does not consider the agency of the readers themselves that can be implied by the headlines. Appealing to the reader is common in political communication and advertis-

Table 2: An example of a headline annotation. Annotators are offered to assign a level of agency to all mentioned noun phrases on a scale from 1 (an entity without agency in this case) to 5 (an entity solely in full control of the situation). A separate annotation is made for any external entities and influences.

Headline	AI Programme Creates Own Language, Researchers Baffled	
Possible Annotation		Numerical value
AI Programme	High level of agency (4)	0.75
Language	No agency (1)	0
Researchers	Low level of agency (2)	0.25
External entities	Not implied (1)	0

ing, but annotating such messages would require a different toolset. But perhaps the most significant limitation is that the current project is limited only to English-language media. If the perceived agency prediction task is reliably solvable, it should be further considered in the multilingual setting. Considering languages where syntactic animacy is more present than in English would potentially introduce additional difficulties. And, as Findor et al. (2021) showed, perceptions of agency shift significantly through seemingly literal translation because of different etymologies and connotations. Therefore, building a multilingual corpus with annotated agency levels may present a much more challenging task.

## Acknowledgments

This research is supported by the Wallenberg AI, Autonomous Systems and Software Program – Humanities and Society (WASP-HS) and the Marianne and Marcus Wallenberg Foundation.

## References

- Gavin Abercrombie, Amanda Cercas Curry, Tanvi Dinkar, Verena Rieser, and Zeerak Talat. 2023. [Migrates. On Anthropomorphism in Dialogue Systems](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4776–4790, Singapore. Association for Computational Linguistics.
- Myra Cheng, Kristina Gligoric, Tiziano Piccardi, and Dan Jurafsky. 2024. [AnthroScore: A Computational Linguistic Measure of Anthropomorphism](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 807–825, St. Julian’s, Malta. Association for Computational Linguistics.
- Paul Engelmann, Peter Trolle, and Christian Hardmeier. 2024. [A Dataset for the Detection of Dehumanizing Language](#). In *Proceedings of the Fourth Workshop on Language Technology for Equality, Diversity, Inclusion*, pages 14–20, St. Julian’s, Malta. Association for Computational Linguistics.
- Andrej Findor, Matej Hruška, Petra Jankovská, and Michaela Pobudová. 2021. [Re-examining public opinion preferences for migrant categorizations: “refugees” are evaluated more negatively than “migrants” and “foreigners” related to participants’ direct, extended, and mass-mediated intergroup contact experiences](#). *International Journal of Intercultural Relations*, 80:262–273.
- Nanna Inie, Stefania Druga, Peter Zukerman, and Emily M. Bender. 2024. [From “AI” to Probabilistic Automation: How Does Anthropomorphization of Technical Systems Descriptions Influence Trust?](#) In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’24*, pages 2322–2347, New York, NY, USA. Association for Computing Machinery.
- Steve Kirkwood. 2017. [The humanisation of refugees: A discourse analysis of UK parliamentary debates on the European refugee ‘crisis’](#). *Journal of Community & Applied Social Psychology*, 27(2):115–125.
- Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. [Detecting Frames in News Headlines and Its Application to Analyzing News Framing Trends Surrounding U.S. Gun Violence](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514, Hong Kong, China. Association for Computational Linguistics.
- Yinhan Liu. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, and Malvina Nissim. 2022a. [Dead or Murdered? Predicting Responsibility Perception in Femicide News Reports](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1078–1090, Online only. Association for Computational Linguistics.

- Gosse Minnema, Sara Gemelli, Chiara Zanchi, Tommaso Caselli, and Malvina Nissim. 2022b. [SocioFillmore: A Tool for Discovering Perspectives](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 240–250, Dublin, Ireland. Association for Computational Linguistics.
- Igor Ryazanov and Johanna Björklund. 2023. How does the language of ‘threat’ vary across news domains? a semi-supervised pipeline for understanding narrative components in news contexts. *Swedish Artificial Intelligence Society*, pages 94–99.
- Igor Ryazanov, Carl Öhman, and Johanna Björklund. 2024. [How ChatGPT Changed the Media’s Narratives on AI: A Semi-Automated Narrative Analysis Through Frame Semantics](#). *Preprint*, arXiv:2408.06120.
- Tazreena Sajjad. 2018. What’s in a name? ‘Refugees’, ‘migrants’ and the politics of labelling. *Race & Class*, 60(2):40–62.
- Isidora Tourni, Lei Guo, Taufiq Husada Daryanto, Fabian Zhafransyah, Edward Edberg Halim, Mona Jalal, Boqi Chen, Sha Lai, Hengchang Hu, Margrit Betke, Prakash Ishwar, and Derry Tanti Wijaya. 2021. [Detecting Frames in News Headlines and Lead Images in U.S. Gun Violence Coverage](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4037–4050, Punta Cana, Dominican Republic. Association for Computational Linguistics.