

Designing NLP Systems That Adapt to Diverse Worldviews

Claudiu Creangă^{2,3}, Liviu P. Dinu^{1,3}

¹ Faculty of Mathematics and Computer Science

² Interdisciplinary School of Doctoral Studies, ³ HLT Research Center

University of Bucharest, Romania

claudiu.creanga@s.unibuc.ro, ldinu@fmi.unibuc.ro

Abstract

Natural Language Inference (NLI) is foundational for evaluating language understanding in AI. However, progress has plateaued, with models failing on ambiguous examples and exhibiting poor generalization. We argue that this stems from disregarding the subjective nature of meaning, which is intrinsically tied to an individual's *weltanschauung* (which roughly translates to worldview). Existing NLP datasets often obscure this by aggregating labels or filtering out disagreement. We propose a perspectivist approach: building datasets that capture annotator demographics, values, and justifications for their labels. Such datasets would explicitly model diverse worldviews. Our initial experiments with a subset of the SBIC dataset demonstrate that even limited annotator metadata can improve model performance.

Keywords: weltanschauung, perspectivism, alignment

1. Introduction

Natural Language Inference (NLI) lies at the heart of developing and evaluating language understanding in AI models. As Montague stated, entailment is "the basic aim of semantics" (Montague, 1970) and a lot of focus has been put in building models that score high on NLI datasets. Recently, two big issues emerged. Firstly, the research has reached a plateau on these datasets where the models perform almost as well as humans on samples where there is high human agreement, but perform poorly on samples with high entropy (from 0.9 to 0.5 accuracy) (Nie et al., 2020a). Secondly, models have poor generalisation abilities and their high score on in-distribution samples doesn't translate to a high score on out-of-distribution samples (Bras et al., 2020), (Zhou and Bansal, 2020), a sign that they actually use shallow heuristics rather than understanding. We will argue that these models lack a worldview and that using perspectivist approaches we can improve our solutions to both problems. A significant obstacle in this research is the scarcity of datasets that preserve annotator demographics and reveal their socio-political values, which are crucial for understanding their worldviews. Our code is open source and available to use on [Anonymous GitHub](#). Our present work focuses on worldviews which are shaped culture, values, beliefs and less by demographic data. Several studies (Orlikowski et al., 2023) have shown that demographic data alone is not a good predictor of annotator's views.

2. Related Work

Basile et al. (2021) offer a valuable synthesis of prior research in perspectivist machine learning,

clearly delineating two distinct approaches within the field: weak and strong perspectivism. Historically, the illusion of ground truth was ingrained in every NLP dataset. At start, NLP datasets were built only with a single annotator per label. That label was taken as the truth, even if the language was ambiguous or the annotator made a mistake. Then, a step forward was made when weak perspectivist research acknowledged the potential for disagreement and errors. It therefore adopted a multi-annotator approach to capture diverse perspectives. However, where there was human disagreement, this was solved either by aggregation (majority voting) or by filtering (removing low agreement samples). But, by removing the low agreement samples, we remove a big part of human speech which is by its virtue ambiguous and, by aggregation, we obscure the rich diversity of valid interpretations inherent in human communication. Strong perspectivist research embraces linguistic diversity, recognizing that even when aiming for a single target label, models benefit from exposure to non-aggregated data (manifesto¹). In the literature we've found three main ways to embrace variation:

- **Multilabel categorical classification:** where a single model predicts multiple labels for each sample (Ferracane et al., 2021), (Jiang and de Marneffe, 2022) and others;
- **Soft label classification:** where a single model is trained on the distribution of labels for a given sample and the model predicts that distribution (Peterson et al., 2019), (Uma, 2020) and others.
- **Radical perspectivist:** where we train a model for each annotator so that a model

¹<https://pdai.info/>

learns the behaviour of one particular annotator (Akhtar et al., 2021). This requires the identification of each annotator, which few datasets have.

Our research continues in the line of the radical perspectivist approach and takes it one step further, by including not only the demographics of the annotators, but also the values and beliefs which make the annotator's worldview.

3. The goal of an NLP system

NLP is a large field and uses of NLP systems can vary from translations to information extraction, summarization and others. But at its core, we can say that an NLP system aims to enable computers to understand and generate human language in a meaningful way. We argue that meaning is subjective and cannot be separated from a worldview. Our language influences how we perceive and understand the world and the other way around. The subjective nature of meaning doesn't entail that each annotator position is equally valid under the current paradigm (an extreme relativistic view that can have negative implications in some tasks like hate speech detection (Curry et al., 2024)), but that each position is part of a worldview that should be dealt with, not erased.

3.1. Meaning and Worldview

Quine's radical indeterminacy (Quine, 1980) shows that no sentence has only one meaning and there is no way to determine the one correct translation of a sentence in another language. This happens because meaning is embedded within the speaker's entire web of beliefs, culture, and how they experience the world. The term that better describes this concept is **Weltanschauung**, roughly translated by worldview. The main problem of NLP systems today is that it doesn't take into account the different worldviews in which a sentence can be interpreted. We cannot know the right label for large category of utterances if we don't contextualise it in a worldview. If we accept Quine's position that there's no such thing as **purely mental meanings** (mental dictionaries from which we take the definition of every word we use) and that what words mean is inextricably linked to how speakers behave and look at the world (*weltanschauung*), then we can agree that the legacy NLP datasets do not offer the information we need. By **legacy datasets** we mean datasets that only have aggregated labels. But non-aggregated labels are not enough and we should go further and include those that do not offer labels by annotator id, demographic metadata and worldview metadata about each annotator.

Here we need to recognize that while many inferences are influenced by worldview, certain types – such as those grounded in formal logic or mathematical reasoning – hold regardless of the annotator (**deductive inferences**). If we were to use logic to establish what inference means, we would denote it by this formula:

$$\forall w \in W : (P(w) \rightarrow H(w))$$

Which means that for every world w , if the premise P is true in world w , then the hypothesis H is also true in world w . Only deductive inferences can pass this type of rigour:

Premise: All men are mortal and Socrate is a man.

Hypothesis: Socrate is mortal.

Label: Entailment.

Here, no matter what we mean by Socrate and men, the inference is still valid. Compared to deductive inferences, **inductive inferences** need to pass a lower bar. Traditionally the standard was what a "common man" would assume to be true or false about an utterance. The creators of datasets sometimes give annotators instructions on how to evaluate utterances (Bowman et al., 2015). In case of SNLI, entailment meant a "definitely true description" and neutral "might be a true description". In case of MNLI the instructions for entailment was "definitely correct" and "might be correct" for neutral. In Gubelmann 2024 the threshold is much more lax, a "good reason" is sufficient for entailment and it is given the following example:

Premise: The streets are wet.

Hypothesis: It has rained.

Label: Entailment.

Firstly, we argue that even a simple inference like this cannot be made without considering the annotator's context. For instance, an annotator living in a town where streets are washed every morning might attribute wet streets to cleaning rather than rain as his first thought. Secondly, providing annotators with instructions biases the results, as it attempts to override their individual understanding of entailment and contradiction. This prevents the study from authentically reflecting common language use which comes in different varieties. And NLP systems should reflect how humans naturally speak.

3.2. Building a worldview-annotated dataset

Consider the following basic pair of sentences:

Premise: It is dark outside.

Hypothesis: It is dangerous to go outside.

There is no right label for this pair and the annotator would use his life experiences to annotate it. A woman would be more likely to annotate it as an entailment. Many men who live in high crime areas would potentially do the same, while others would see no good reason to be dangerous outside if it is dark. Recently built perspectivist datasets, such as ChaosNLI (Nie et al., 2020b) which collects 100 annotations for each label, but then proceeds to remove the worldviews of annotators and provide only a distribution (i.e. 30% E, 50% N, 20% C) make the same mistake of weak perspectivist research. Relying solely on a distribution, one cannot learn in which worldview this statement is entailment or not and if we don't know the annotator backgrounds and how diverse they are, we risk to capture only a limited range of potential interpretations.

In Figure 1 we explain how we can build worldview-annotated datasets. It is necessary to have a diverse pool of annotators aligned to the task. For instance, a BioNLP task would benefit from annotators with expertise in health sciences. Metadata should be collected about each of these annotators: demographics and values. Worldviews are part of a paradigm, a set of principles shared by every annotator at a certain point in time. Each annotator should label items according to their worldview, while being mindful of potential "noise" – errors caused by factors like lack of attention, which are unrelated to their perspective. To mitigate noise and preserve valid interpretations, annotators should provide justifications for their labels. They will then self-review their labels based on their own explanation and reconsider the label or not. This self-review step should remove some noise. The **justifications** are important for a second reason. How can we address the scenario where two annotators assign the same label, but their justifications diverge so significantly that the apparent agreement is misleading? This case suggests that reasons must be kept in the dataset and if reasons diverge significantly there should be different categories in each label. A model that learns a worldview should arrive at a label through the same reasoning as the annotator. Unfortunately there is no public dataset where reasons are available, but we aim to build one and we hypothesize that it would help the models in learning a worldview.

Training a model in this way means modelling a worldview in which the model doesn't pretend universality and is by essence local. We hypothesize that this grounding in a worldview will make

the model more consistent in its judgements and generalise better.

4. Modelling a worldview

Unfortunately, the demographic metadata of annotators is removed in most datasets and there are very few datasets that collected annotators' social and political values in order to create a worldview. However, this is changing, with at least four datasets that we are aware of now including more information: Chulvi 2023, Sachdeva et al. 2022, Sap et al. 2019 and Hettiachchi 2023. In (Sap et al., 2022) we see that annotators with stronger racist beliefs demonstrate a tendency to mislabel African American English (AAE) as toxic, while being less likely to identify anti-Black language as harmful, suggesting that even only demographic metadata can be helpful. To the best of our knowledge, SBIC (Sap et al., 2019) is the most complete dataset in regards to including both annotator demographics and the political beliefs of the annotator (liberal or conservative). Their annotators are from U.S. and Canada and, although they are gender and age-balanced, the ethnicity is not diverse, there are a lot more white (82%) than any other group (4% Asian, 4% Hispanic, 4% Black). They annotated posts from Reddit and Twitter for offensiveness, intent to offend, sexual references and, if it was offensive, which groups did it target.

We used this dataset for our goal of building models with demographic and political metadata. Because the dataset was built with detecting Social Bias Frames in mind, they had 263 different annotators and at most 3 annotators per post. For our use case we had to find annotators from different backgrounds that annotated the same posts. The biggest subset that we could find from this dataset was for 2 annotators, man (liberal, white) and woman (mod-conservative, white), which together annotated 290 samples (worker ids are hidden due privacy concerns).

Given it is a small dataset, we are aware of the limitations of our results. We used K-Fold validation (10) to reduce the risk of overfitting and we used a pre-trained DeBERTaV3 model (He et al., 2021) followed by a fully connected layer for predictions. We used a learning rate of 5e-5 for 3 epochs where we trained only the output layer and then a small learning rate of 2e-5 for one epoch where we trained the output layer and the last layer from the pre-trained model. The target label is if the post is offensive or not.

As shown in (Table 1), leveraging annotator metadata provides a boost in performance. Training on aggregated data without metadata yields a test F1 score of 0.3. While splitting the data by annotator, training 2 separated models and ag-

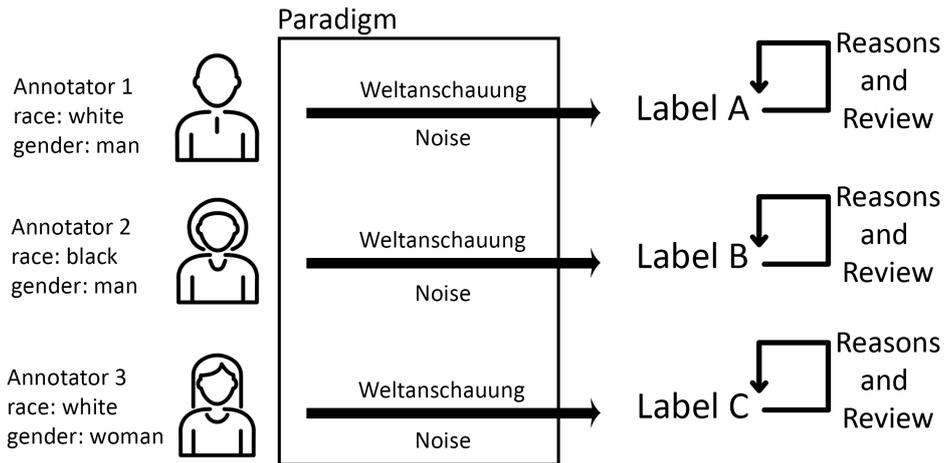


Figure 1: Building a worldview-annotated dataset. It is necessary to have a diverse pool of annotators aligned to the task. Metadata should be collected about each of these annotators: demographic and values. Each annotator should label items according to their worldview, while being mindful of potential "noise" – errors caused by factors like inattention, which are unrelated to their perspective. To mitigate noise and preserve valid interpretations, annotators should provide justifications for their labels. They will then self-review their labels based on their own explanation.

Table 1: F1 scores of the three types of training: training on aggregated data, using 2 models for each annotator with and without metadata.

Training Type	Val score	Test score
Aggregated data	0.3	0.3
2 models no metadata	0.37	0.36
2 models and metadata	0.38	0.38

gregating the outputs afterwards improves the F1 score to 0.36. Instead, if we also concatenate the annotators' metadata to the input before the encoding layers, we increase the F1 score to 0.38. Even though the dataset is small, there is an improvement when we consider the annotators' background.

5. Conclusion

Traditional NLP approaches have exhibited limitations in both generalization and performance on challenging examples within NLI datasets. Part of these shortcomings stem from the omission of the subjective nature of meaning and the diverse worldviews that shape individual interpretations of language. For inductive inferences, the notion of a universally "correct" label detached from a worldview is misleading.

We propose that embracing perspectivist approaches and building worldview-annotated datasets is crucial for advancement in NLP. Such datasets must capture annotator justifications, demographic information, and the values that com-

prise their worldview, offering a richer understanding of linguistic variation. Such datasets are missing at the moment, but initial experiments with a subset of the SBIC data support our hypothesis: even with limited demographics and a basic political orientation label, incorporating annotator metadata improves model performance.

Acknowledgements

This work was partially supported by a grant on Machine Reading Comprehension from Accenture Labs and by the POCIDIF project in Action 1.2. "Romanian Hub for Artificial Intelligence".

Sohail Akhtar, Valerio Basile, and Viviana Patti. 2021. [Whose opinions matter? perspective-aware models to identify opinions of hate speech victims in abusive language detection.](#)

Valerio Basile, Federico Cabitza, Andrea Campagner, and Michael Fell. 2021. [Toward a perspectivist turn in ground truthing for predictive computing.](#)

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference.](#) In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.

- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew E. Peters, Ashish Sabharwal, and Yejin Choi. 2020. [Adversarial filters of dataset biases](#).
- Fontanella-L. Labadie-Tamayo R. Rosso P. Chulvi, B. 2023. [Social or individual disagreement? perspectivism in the annotation of sexist jokes](#).
- Amanda Cercas Curry, Gavin Abercrombie, and Zeerak Talat. 2024. [Subjective *Isms*? on the danger of conflating hate and offence in abusive language detection](#).
- Elisa Ferracane, Greg Durrett, Junyi Jessy Li, and Katrin Erk. 2021. [Did they answer? subjective acts and intents in conversational discourse](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1626–1644, Online. Association for Computational Linguistics.
- Katis-I. Niklaus-C. et al Gubelmann, R. 2024. [Capturing the varieties of natural language inference: A systematic survey of existing datasets and two novel benchmarks](#). page 21–48.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Holcombe-James I.-Livingstone S. de Silva A. Lease M. Salim F. D. Sanderson M. Hettiachchi, D. 2023. [How crowd worker factors influence subjective annotations: A study of tagging misogynistic hate speech in tweets](#).
- Nan-Jiang Jiang and Marie-Catherine de Marneffe. 2022. [Investigating reasons for disagreement in natural language inference](#). *Transactions of the Association for Computational Linguistics*, 10:1357–1374.
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020a. [What can we learn from collective human opinions on natural language inference data?](#)
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020b. [What can we learn from collective human opinions on natural language inference data?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Matthias Orlikowski, Paul Röttger, Philipp Cimiano, and Dirk Hovy. 2023. [The ecological fallacy in annotation: Modelling human label variation goes beyond sociodemographics](#).
- Joshua C. Peterson, Ruairidh M. Battleday, Thomas L. Griffiths, and Olga Russakovsky. 2019. [Human uncertainty makes classification more robust](#).
- W. V. O. Quine. 1980. Two dogmas of empiricism. from a logical point of view. page 20–46.
- Pratik Sachdeva, Renata Barreto, Geoff Bacon, Alexander Sahn, Claudia von Vacano, and Chris Kennedy. 2022. [The measuring hate speech corpus: Leveraging rasch measurement theory for data perspectivism](#). In *Proceedings of the 1st Workshop on Perspectivist Approaches to NLP @LREC2022*, pages 83–94, Marseille, France. European Language Resources Association.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2019. [Social bias frames: Reasoning about social and power implications of language](#).
- Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A. Smith. 2022. [Annotators with attitudes: How annotator beliefs and identities bias toxic language detection](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5884–5906, Seattle, United States. Association for Computational Linguistics.
- Fornaciari T. Hovy-D. Paun S. Plank B. Poesio M. Uma, A. 2020. [A case for soft loss functions](#).
- Xiang Zhou and Mohit Bansal. 2020. [Towards robustifying NLI models against lexical dataset biases](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8759–8771, Online. Association for Computational Linguistics.