

# Position: From Noise to Signal to Selbstzweck - Reframing Human Label Variation in the Era of Post-training in NLP

Shanshan Xu<sup>1,2</sup>, Santosh T.Y.S.S<sup>3</sup>, Barbara Plank<sup>4</sup>

<sup>1</sup>Department of Computer Science, University of Copenhagen, Denmark

<sup>2</sup>Faculty of Law, University of Copenhagen, Denmark

<sup>3</sup>Amazon <sup>4</sup>LMU Munich & Munich Center for Machine Learning (MCML)

shanshan.xu@di.ku.dk, santoshtyss@gmail.com, b.plank@lmu.de

## Abstract

Human Label Variation (HLV) refers to legitimate disagreement in annotation that reflects the diversity of human perspectives rather than mere error. Long treated in NLP as *noise* to be eliminated, HLV has only recently been reframed as a *signal* for improving model robustness. With the rise of large language models (LLMs) and post-training methods such as human feedback-based alignment, the role of HLV has become increasingly consequential. Yet current preference-learning datasets routinely collapse multiple annotations into a single label, flattening diverse perspectives into artificial consensus. Preserving HLV is necessary not only for pluralistic alignment but also for sociotechnical safety evaluation, where model behavior must be assessed in relation to human interaction and societal context. This position paper argues that preserving HLV as an embodiment of human pluralism must be treated as a *Selbstzweck*, an intrinsic value in itself. We analyze the limitations of existing preference datasets and propose actionable strategies for incorporating HLV into dataset construction to better preserve pluralistic human values.

## 1 Introduction

The belief in a single “gold” human label runs deep in NLP. For decades, annotation pipelines have been designed to minimize inter-annotator disagreement by enforcing a single gold label, treating disagreement as noise to be discarded. Majority vote aggregation, while convenient, risks to flatten the diverse human perspectives. Recently, however, a growing body of work has begun to question this tradition, challenging the assumption that low inter-annotator agreement is necessarily a sign of poor data quality (Plank et al., 2014; Fornaciari et al., 2021). Annotation disagreement is increasingly observed in modern NLP tasks, especially in tasks involving social context.

Growing work in NLP acknowledges that disagreement is not merely an artifact of error, but often an unavoidable and legitimate (Aroyo and Welty, 2015; Basile et al., 2021; Prabhakaran et al., 2021). This stands in contrast to clear-cut annotation errors, such as misreading task instructions or misclicking a label. To capture this crucial distinction, Plank 2022 introduced the notion of human label variation (HLV), urging the community to recognize HLV as a meaningful signal that reflects the pluralistic nature of human values.

HLV is not only a feature of annotation but also a key channel through which pluralistic social perspectives enter technical systems. Research on AI safety highlights that AI systems are sociotechnical, integrating both model and human components (Weidinger et al., 2023). Evaluations of AI must go beyond assessing model capabilities alone and also account for broader social contexts (Solaiman et al., 2023; Weidinger et al., 2024). For example in tasks such as hate speech detection, this entails accounting for annotators social identities and the societal structures that shape definitions of harm.

Despite growing awareness of pluralistic human alignment (Parappan and Henao, 2025; Sorensen et al., 2025), most efforts focus on evaluation frameworks (Sorensen et al., 2024) or reward-learning algorithms (Chen et al., 2025), while the foundation of alignment, the datasets encoding human preferences remains underexamined. Zhang et al. 2025 investigates annotator disagreement in preference dataset (Fig 2 in App § 2), revealing that annotators’ diverging preference is systematic rather than random noise, and that aggregation can undermine pluralistic alignment. While their study provides valuable empirical insights, our paper takes a broader conceptual and prescriptive perspective.

Our contributions are threefold. First, we propose a three-stage conceptual framework (Fig 1) tracing the conceptualization shifts of HLV from *noise* to be eliminated, to *signal* to be modeled and

Stage	Human Adjacent	Human Aware	Human Centered
Tasks	Expert systems, NER, POS, etc (Ashley, 1991)	Everyday language use: Sentimental Analysis, hate speech detection, etc (Pérez-Rosas et al., 2018)	Interactive and creative open-ended tasks: pair programming, creative writing, counseling (Litschko et al., 2023)
Evaluation	<b>Capability:</b> technical components in isolation (Jones and Galliers, 1995)	<b>Human Interaction:</b> bias, fairness and interpretability (Belinkov et al., 2020; Mehrabi et al., 2021)	<b>Systemic impact:</b> broader impacts on the social systems (Weidinger et al., 2023)
Role of Humans	<b>Peripheral:</b> largely outside the system's	<b>Acknowledged:</b> humans as data annotator and end-users	<b>Central:</b> humans as co-creators
View of HLV	<b>Noise</b> to be discarded	<b>Signal</b> to improve performance	<b>Selbstzweck</b> to align diverse human values

Figure 1: Three stages of NLPs evolution toward human-centered systems. The figure illustrates how the changing view of HLV parallels the the shift in evaluation practices, and growing integration of human pluralism in NLP research. Evaluation row integrates the sociotechnical perspectives of Weidinger et al. 2023.

to *Selbstzweck* as intrinsic value. We also explain how these shifts parallel the integration of human pluralism into NLP technologies, as well as the evaluation focus shift in sociotechnical framework<sup>1</sup> (§ 3). Second, we critically examine current preference datasets and identify key limitations that prevent them from capturing genuine human pluralism (§ 2). Third, we provide actionable strategies for embedding HLV in future alignment pipelines<sup>2</sup> (§ 4). Rather than presenting new empirical results, we synthesize existing evidence into a coherent framework that reframes how the community should approach human variation in the post-training era.

## 2 Human Label Variation

A growing body of work has studied annotation disagreement across NLP tasks, proposing taxonomies to identify its potential sources (Uma et al., 2021; Sandri et al., 2023; Jiang and Marneffe, 2022; Xu et al., 2024a). Recent work (Weber-Genzel et al., 2024a) emphasizes the importance of distinguishing *annotation errors* from *HLV*. Annotation errors arise when labels are assigned for invalid reasons. For example, annotator sloppiness, misunderstanding of instructions, or technical issues with the interface. These can often be reduced through better training, clearer guidelines, or improved UI design. By contrast, HLV reflects genuine task ambiguity and human subjectivity, which cannot and should not be eliminated.

As NLP has progressed from well-defined tasks

<sup>1</sup>App A includes the introduction of the sociotechnical evaluation framework (Weidinger et al., 2023).

<sup>2</sup>HLV impacts full ML pipeline: data, modeling and evaluation (Plank, 2022). This 4-page position paper focuses on datasets as a key starting point for pluralistic alignment. Further, we focus on post-training data, as it explicitly collects human judgments where annotator-level variation can be directly observed, whereas pretraining corpora consist of passively gathered text.

(e.g., POS tagging, NER) toward socially embedded and subjective problems such as emotion categorisation (Demszky et al., 2020) and humorousness prediction (Simpson et al., 2019), HLV has become increasingly prevalent and unavoidable. Even in natural language inference (NLI), which once considered an objective benchmark, researchers observe systematic disagreement (See examples in App B) (Pavlick and Kwiatkowski, 2019; Jiang et al., 2023). HLV is also found in preference datasets, as shown in Fig 2 in App C. Zhang et al. (2025) shows that over 30% of examples in Multi-Pref (Miranda et al., 2025) and HelpSteer2 (Wang et al., 2024) contain disagreement, and that fewer than 25% of these cases should be attributed to pure annotation error. Nevertheless, major public preference datasets (Ouyang et al., 2022; Bai et al., 2022) are either annotated by a single annotator or aggregated into a single consensus label. Treating all disagreement as error risks erasing HLV and thereby flattening the pluralism in human values that these datasets could otherwise capture.

## 3 Conceptual Evolution of HLV

As NLP tasks have grown increasingly subjective, the conceptualisation of HLV has shifted: from being discarded as *noise*, to being recognized as a valuable *signal*. We argue that a further conceptual shift is required in the era of LLMs: HLV, as an embodiment of pluralistic human values, should be elevated to *Selbstzweck*, a guiding *design principle* for building genuinely human-centered AI. This reconceptualization reflects both the increasing subjectivity of NLP tasks and the broader re-positioning of human pluralism within NLP system design. Winter (2024) propose a framework that categorizes the role of humans in relation to technology, highlighting different levels of integration of human values and agency. Building on this perspective,

we introduce a three-stage framework (Fig 1). The trajectory from *Human Adjacent* → *Human Aware* → *Human Centered* captures how human pluralism have been positioned in NLP over time, moved from the periphery to the core of system design.

**Human Adjacent:** In the early era of NLP (before the 1990s), humans occupied a largely peripheral position in NLP systems. Systems were designed by research labs and corporations primarily to optimize narrowly defined tasks such as text retrieval, syntactic parsing, speech recognition, or named entity recognition. Most systems relied on expert knowledge encoded in handcrafted rules or symbolic networks, e.g., HYPO (Ashley, 1991) for legal reasoning in U.S. trade law. The dominant motivation was to increase organizational efficiency and, in the corporate setting, profitability. These tasks were assumed to have a single correct solution, annotated by domain experts. Annotation disagreement was conceptualized as mere error, to be eliminated in order to preserve the gold standard. Evaluation focus was on the *capability level*, testing the technical components in isolation, and centering on accuracy and F1 scores against presumed objective truth (Jones and Galliers, 1995). In this stage, humans were *adjacent*: ordinary people largely excluded from shaping how systems were designed.

**Human Aware:** From the 1990s onward, with the rise of the web, deep learning, and later transformer-based models such as BERT (Devlin et al., 2019), NLP shifted from specialized organizational tools to applications embedded in everyday life, including fake news detection (Pérez-Rosas et al., 2018) and hate speech detection (Sap et al., 2019). Annotation practices also changed: large-scale crowdsourcing replaced reliance on domain experts, making it possible to scale data collection but also surfacing greater HLV across annotators. At the same time, researchers began to recognize that NLP applications, especially those used in high-impact domains, can not be tested in isolation without social context. Models’ outputs can affect peoples opportunities or safety in applications like HR screening (Li et al., 2020) and legal decision support (Xu et al., 2023a). Beyond task technical capability, evaluation focus shifted toward *Human-interaction level*, with growing research on interpretability (Mehrabi et al., 2021), bias and fairness (Belinkov et al., 2020). In this stage, NLP became *human aware*: systems were still largely

designed and deployed by organizations, but they increasingly targeted ordinary people as end-users in everyday contexts such as social media.

**Human centered:** Since the launch of ChatGPT (OpenAI, 2023), LLMs have transformed NLP from a domain of specialized, compartmentalized tasks (Litschko et al., 2023) into one of open-ended, interactive systems. As users increasingly turn to these models for factual information and practical advice, their societal influence has grown rapidly. Studies demonstrate that LLMs can shape users decisions and political beliefs (Potter et al., 2024). From a sociotechnical perspective, evaluating LLMs requires more than assessing model capabilities or interactions. It demands attention to their *systemic impact* (Weidinger et al., 2023): the broader effects on social, economic, and environmental systems, including consequences that appear only at scale, such as political polarization or shifts in public trust. These developments make the question of alignment more urgent than ever. To move toward genuinely *human-centered* AI, pluralism must be treated as a *Selbstzweck*, a guiding design principle. Incorporating HLV proactively in data collection, modelling and evaluation is essential for building systems that are not only aligned but also adaptive to the diversity of human needs, contexts, and worldviews.

## 4 Integrating HLV in Preference Dataset

**Open Issues** In the era of LLM, learning from human preferences has emerged as the de facto post-training method for AI Alignment. The purpose of a preference dataset is to capture human judgments about which model outputs are better, more helpful and harmless. Such judgments are inherently subjective. Annotators may differ in how they weigh factual accuracy, conciseness, empathy, or even stylistic tone, as illustrated in Fig 2. These differences are not annotation errors but expressions of HLV, reflecting the diversity and pluralism of human values. The current approach of discarding all the disagreements as mere annotation error<sup>3</sup> risks over-reducing the inherent HLV which captures the pluralism in human values.

<sup>3</sup>Distinguishing genuine HLV from annotation errors remains a practical challenge; we provide further discussion in App E.

## 4.1 Our Suggestions

**Carefully Select the Annotator Pool** To promote pluralism in human values, managing subjectivity is not about eliminating HLV but about shaping the annotator pool so that it reflects the diversity of the target population. Building preferences datasets is analogous to conducting a survey: a narrow or opaque selection of annotators risks amplifying bias in downstream models (Eckman et al., 2024). Yet most LLMs disclose little about annotator recruitment or training; recent analyses (Chalkidis, 2025) show that after 2022, major models virtually stopped reporting on annotator backgrounds or training practices. To mitigate the selective bias in selecting annotator pool, survey methodology offers guidance, for instance: define the target population, recruit with stratified quotas, and use multi-channel strategies to include underrepresented groups (Ahmed, 2024). Collecting annotator metadata further enables balanced assignments and post-hoc corrections such as post-stratification weighting. For example, Aroyo et al. 2023 carefully hired a wide range of annotators to collect diverse opinions of AI safety, providing a basis for analyzing how ratings intersect with demographic categories.

**Release Annotator-Level Preferences** The original HelpSteer2 dataset (Wang et al., 2024) provided only mean scores from five annotators, but its recent update includes the individual annotations. We call for broader adoption of this practice, echoing earlier suggestions on data annotation (Prabhakaran et al., 2021). Annotator-level labels enable distributional or mixture reward models that capture uncertainty and heterogeneity in preferences rather than collapsing them into a single scalar, leading to more representative policies. Zhang et al. (2025) further show that aggregating labels during preference training can harm pluralistic alignment: LLMs are rewarded equally for making decisive choices in both high- and low-agreement cases. This tendency also manifests in LLM-as-Judge evaluation pipelines, which often force a single winning response even when preferences diverge. Further, the fine-grained annotations would also allow investigation on the *tied preferences*, where annotators genuinely express no preference between responses (see App D for further discussion on HLV vs. tied preferences). Finally, releasing unaggregated annotator-level preferences enhances transparency and enables future research

into the sources of disagreement, an essential step toward the pluralistic human alignment.

**Leverage Diverse Types of Feedback** Current preference datasets typically rely on two feedback protocols: relative (pairwise comparisons; Wirth et al. 2017) and absolute (e.g., Likert-scale ratings; Arzate Cruz and Igarashi 2020). While relative feedback is simple and widely adopted, it tends to favor assertive or rhetorically salient responses (Kaufmann et al., 2023), obscuring nuanced diverse human judgments. Even with annotator-level labels, binary pairwise comparisons still collapse nuanced diversity in human preferences, whereas absolute feedback better preserve fine-grained perspectives (Zhang et al., 2025). Recent studies further demonstrate that combining diverse feedback types, including ratings, comparisons and descriptive evaluations, can produce reward models and downstream policies that surpass single-feedback baselines (Metz et al., 2025). To address these limitations, we suggest to leverage heterogeneous feedback protocols (e.g., scalar ratings, pairwise and n-wise comparisons and demonstrations), to faithfully capture the pluralism of human values that alignment aims to preserve<sup>4</sup>.

**Pluralism as Reasoning** In many real-world settings, a single operative decision is unavoidable, for instance, a court must ultimately issue one binding ruling. In such settings, AI systems can preserve pluralism in the reasoning process rather than the final outcome: instead of collapsing disagreement into a single label, systems can document, aggregate, and justify divergent perspectives alongside the unified decision. Legal practice offers a useful analogue: when the judges can not reach a unanimous decision, courts issue binding rulings by majority vote while formally recording dissenting opinions, preserving diverse perspectives for accountability and for future legal reasoning (Xu et al., 2024b). Another possible approach is AI as a mediator. Recent work shows an LLM mediator (Tessler et al., 2024) can help small groups to find common ground on divisive political issues by synthesizing participants written views into shared group statements.

---

<sup>4</sup>Feedback design and alignment algorithms are interconnected. App§ 5 provides ideas for alignment algorithms

## 5 Conclusion

In this paper, we trace the evolving view of Human Label Variation (HLV) from *noise*, to *signal*, to *Selbstzweck*; and illustrate how this shift parallels the broader integration of human pluralism into NLP technologies. To advance toward genuinely human-centered AI, we argue for the proactive inclusion of HLV in the creation of preference datasets. While our discussion centers on preference data due to scope limit, the implications of HLV extend across the entire machine learning pipeline, from data collection, to model training, and evaluation. We hope this survey raises awareness within the community and inspires further work on modeling and evaluating alignment in ways that account for pluralistic human values.

### Limitations

This position paper aims to be concise while synthesizing a broad conceptual shift: the evolving view of Human Label Variation (HLV) mirrors the broader integration of human pluralism into NLP technologies. As such, several limitations should be acknowledged:

First, we acknowledge that incorporating pluralism introduces additional computational and financial costs. Our goal is not to prescribe a specific implementation strategy, but to emphasize that HLV and pluralistic alignment should be treated as first-order design considerations. The three-stage evolution we outline (Human Adjacent → Human Aware → Human Centered) reflects precisely how technological and societal developments have made it both feasible and necessary to evaluate models beyond pure task performance, even at additional cost.

Second, our suggestions on how to including HLV in dataset creation focuses on preference dataset. Limited by the scope, we do not address the *pretraining* corpora, which likely exerts an even stronger influence on models ideological tendencies and value alignment (Xu et al., 2025; Ceron et al., 2025). The *Selbstzweck* principle may generalize to pretraining data collection, e.g., through normative considerations in data filtering or multilingual corpus construction. We focus on post-training datasets, as post-training data like RLHF explicitly collects human judgments where annotator-level variation can be directly observed, whereas pretraining corpora consist of passively gathered text.

Third, our analysis does not comprehensively cover the whole pipeline of alignment techniques. As pointed out by Plank 2022, implications of HLV span the entire machine learning pipeline, from data collection and model training to evaluation. While we have addressed on the preference dataset collection in § 4 and briefly discussed on alignment algorithms in § 5, future work could explore new evaluation metrics, for example by developing metrics such as:

- **Pluralism fidelity** assesses how closely the models conditional output distribution matches the empirical human response distribution for a given context. This is related to distributional pluralism in recent work (Sorensen et al., 2023) on pluralistic alignment roadmap. In practice, pluralism fidelity can be quantified using standard divergence measures such as KL divergence, or Wasserstein distance as used in prior work (Durmus et al., 2023; Santurkar et al., 2023).
- **Minority regret** quantifies the utility loss (regret) for minority preference groups relative to an optimal group-aligned model, as defined by in Raghavan et al. 2018. This mirrors recent work such as MaxMin-RLHF (Chakraborty et al., 2024), which optimizes for the worstserved group rather than the population average.

Finally, given the rapid pace of progress in this area, some very recent developments may not be included. As a position paper, our aim is to provoke discussion and outline future research directions, rather than to offer comprehensive solutions or empirical evaluations. We encourage further work that operationalizes these principles in a broader range of pluralistic, linguistic, and technological settings. We hope this paper contributes to raising awareness within the community and encourages further exploration of methods for collection, modeling and evaluating human pluralism into NLP research and practice.

### Information about use of AI assistants

In the preparation of this work, the authors utilized ChatGPT and Gemini to polish the writing and improve the coherence of the manuscript. Before the submission, the authors conducted a thorough review and made necessary edits to the content,

taking full responsibility for the final version of the text.

## References

- Sirwan Khalid Ahmed. 2024. How to choose a sampling technique and determine sample size for research: A simplified guide for researchers. *Oral Oncology Reports*, 12:100662.
- Lora Aroyo, Alex Taylor, Mark Diaz, Christopher Homan, Alicia Parrish, Gregory Serapio-García, Vinodkumar Prabhakaran, and Ding Wang. 2023. Dices dataset: Diversity in conversational ai evaluation for safety. *Advances in Neural Information Processing Systems*, 36:53330–53342.
- Lora Aroyo c and Chris Welty. 2015. Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1):15–24.
- Christian Arzate Cruz and Takeo Igarashi. 2020. [A survey on interactive reinforcement learning: Design principles and open challenges](#). In *Proceedings of the 2020 ACM Designing Interactive Systems Conference, DIS '20*, page 11951209, New York, NY, USA. Association for Computing Machinery.
- Kevin D. Ashley. 1991. [Modeling legal argument - reasoning with cases and hypotheticals](#). In *Artificial intelligence and legal reasoning*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1 others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Valerio Basile, Michael Fell, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, Massimo Poesio, and Alexandra Uma. 2021. [We need to consider disagreement in evaluation](#). In *Proceedings of the 1st Workshop on Benchmarking: Past, Present and Future*, pages 15–21, Online. Association for Computational Linguistics.
- Yonatan Belinkov, Sebastian Gehrmann, and Ellie Pavlick. 2020. Interpretability and analysis in neural nlp. In *Proceedings of the 58th annual meeting of the association for computational linguistics: tutorial abstracts*, pages 1–5.
- Tanise Ceron, Dmitry Nikolaev, Dominik Stammach, and Debora Nozza. 2025. What is the political content in llms’ pre-and post-training data? *arXiv preprint arXiv:2509.22367*.
- Souradip Chakraborty, Jiahao Qiu, Hui Yuan, Alec Koppel, Furong Huang, Dinesh Manocha, Amrit Bedi, and Mengdi Wang. 2024. Maxmin-rlhf: Towards equitable alignment of large language models with diverse human preferences. In *ICML 2024 Workshop on Models of Human Feedback for AI Alignment*.
- Ilias Chalkidis. 2025. Decoding alignment: A critical survey of llm development initiatives through value-setting and data-centric lens. *arXiv preprint arXiv:2508.16982*.
- Daiwei Chen, Yi Chen, Aniket Rege, and Ramya Korlakai Vinayak. 2025. Pal: Pluralistic alignment framework for learning from heterogeneous preferences. URL <https://arxiv.org/abs/2406.08469>.
- Jinghong Chen, Guangyu Yang, Weizhe Lin, Jingbiao Mei, and Bill Byrne. 2024. On extending direct preference optimization to accommodate ties. *arXiv preprint arXiv:2409.17431*.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. GoEmotions: A Dataset of Fine-Grained Emotions. In *58th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Esin Durmus, Karina Nguyen, Thomas I Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, and 1 others. 2023. Towards measuring the representation of subjective global opinions in language models. *arXiv preprint arXiv:2306.16388*.
- Stephanie Eckman, Barbara Plank, and Frauke Kreuter. 2024. [Position: Insights from survey methodology can improve training data](#). In *International Conference on Machine Learning*.
- Tommaso Fornaciari, Alexandra Uma, Silviu Paun, Barbara Plank, Dirk Hovy, Massimo Poesio, and 1 others. 2021. Beyond black & white: Leveraging annotator disagreement via soft-label multi-task learning. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- Armin Ghane-Kanafi and Esmail Khorram. 2015. [A new scalarization method for finding the efficient frontier in non-convex multi-objective problems](#). *Applied Mathematical Modelling*, 39:7483–7498.
- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. 2023. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. *arXiv preprint arXiv:2310.11564*.
- 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.

- Nan-Jiang Jiang, Chenhao Tan, and Marie-Catherine de Marneffe. 2023. Understanding and predicting human label variation in natural language inference through explanation. *arXiv preprint arXiv:2304.12443*.
- Karen Sparck Jones and Julia R Galliers. 1995. Evaluating natural language processing systems: An analysis and review.
- Timo Kaufmann, Paul Weng, Viktor Bengs, and Eyke Hüllermeier. 2023. [A survey of reinforcement learning from human feedback](#). *ArXiv*, abs/2312.14925.
- Changmao Li, Elaine Fisher, Rebecca Thomas, Steve Pittard, Vicki Hertzberg, and Jinho D. Choi. 2020. [Competence-level prediction and resume & job description matching using context-aware transformer models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8456–8466, Online. Association for Computational Linguistics.
- Robert Litschko, Max Müller-Eberstein, Rob van der Goot, Leon Weber-Genzel, and Barbara Plank. 2023. [Establishing trustworthiness: Rethinking tasks and model evaluation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 193–203, Singapore. Association for Computational Linguistics.
- Shang Liu, Yu Pan, Guanting Chen, and Xiaocheng Li. 2024. Reward modeling with ordinal feedback: Wisdom of the crowd. *arXiv preprint arXiv:2411.12843*.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35.
- Yannick Metz, Andras Geiszl, Raphaël Baur, and Menatallah El-Assady. 2025. Reward learning from multiple feedback types. In *The Thirteenth International Conference on Learning Representations*.
- Lester James Validad Miranda, Yizhong Wang, Yanai Elazar, Sachin Kumar, Valentina Pyatkin, Faeze Brahman, Noah A. Smith, Hannaneh Hajishirzi, and Pradeep Dasigi. 2025. [Hybrid preferences: Learning to route instances for human vs. AI feedback](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7162–7200, Vienna, Austria. Association for Computational Linguistics.
- Yixin Nie, Xiang Zhou, and Mohit Bansal. 2020. [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)*, pages 9131–9143, Online. Association for Computational Linguistics.
- OpenAI. 2023. [Gpt-4 technical report](#). Accessed: 2023-09-27.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Mohammed Fayiz Parappan and Ricardo Henao. 2025. [Learning subjective label distributions via sociocultural descriptors](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20322–20338, Suzhou, China. Association for Computational Linguistics.
- Ellie Pavlick and Tom Kwiatkowski. 2019. [Inherent disagreements in human textual inferences](#). *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2018. [Automatic detection of fake news](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3391–3401, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Barbara Plank. 2022. The problem of human label variation: On ground truth in data, modeling and evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10671–10682.
- Barbara Plank, Dirk Hovy, and Anders Søgaard. 2014. [Linguistically debatable or just plain wrong?](#) In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 507–511, Baltimore, Maryland. Association for Computational Linguistics.
- Yujin Potter, Shiyang Lai, Junsol Kim, James Evans, and Dawn Song. 2024. [Hidden persuaders: LLMs’ political leaning and their influence on voters](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4244–4275, Miami, Florida, USA. Association for Computational Linguistics.
- Vinodkumar Prabhakaran, Aida Mostafazadeh Davani, and Mark Diaz. 2021. [On releasing annotator-level labels and information in datasets](#). In *Proceedings of the Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 133–138, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Manish Raghavan, Aleksandrs Slivkins, Jennifer Vaughan Wortman, and Zhiwei Steven Wu. 2018. The externalities of exploration and how data diversity helps exploitation. In *Conference on Learning Theory*, pages 1724–1738. PMLR.
- Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa Shukor, Laure Soulier, and Matthieu Cord. 2023. Rewarded

- soups: towards pareto-optimal alignment by interpolating weights fine-tuned on diverse rewards. *Advances in Neural Information Processing Systems*, 36:71095–71134.
- Marta Sandri, Elisa Leonardelli, Sara Tonelli, and Elisabetta Jezeq. 2023. [Why don't you do it right? analysing annotators' disagreement in subjective tasks](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2428–2441, Dubrovnik, Croatia. Association for Computational Linguistics.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. [Whose opinions do language models reflect?](#) In *International Conference on Machine Learning*, pages 29971–30004. PMLR.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The risk of racial bias in hate speech detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.
- Edwin Simpson, Erik-Lân Do Dinh, Tristan Miller, and Iryna Gurevych. 2019. [Predicting humorousness and metaphor novelty with gaussian process preference learning](#). In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 5716–5728.
- Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Canyu Chen, Hal Daumé III, Jesse Dodge, Isabella Duan, and 1 others. 2023. [Evaluating the social impact of generative ai systems in systems and society](#). *arXiv preprint arXiv:2306.05949*.
- Taylor Sorensen, Liwei Jiang, Jena Hwang, Sydney Levine, Valentina Pyatkin, Peter West, Nouha Dziri, Ximing Lu, Kavel Rao, Chandra Bhagavatula, and 1 others. 2023. [Value kaleidoscope: Engaging ai with pluralistic human values, rights, and duties](#). *arXiv preprint arXiv:2309.00779*.
- Taylor Sorensen, Pushkar Mishra, Roma Patel, Michael Henry Tessler, Michiel A. Bakker, Georgina Evans, Iason Gabriel, Noah Goodman, and Verena Rieser. 2025. [Value profiles for encoding human variation](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2047–2095, Suzhou, China. Association for Computational Linguistics.
- Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Miresghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, and 1 others. 2024. [Position: a roadmap to pluralistic alignment](#). In *Proceedings of the 41st International Conference on Machine Learning*, pages 46280–46302.
- Michael Henry Tessler, Michiel A Bakker, Daniel Jarrett, Hannah Sheahan, Martin J Chadwick, Raphael Koster, Georgina Evans, Lucy Campbell-Gillingham, Tantum Collins, David C Parkes, and 1 others. 2024. [Ai can help humans find common ground in democratic deliberation](#). *Science*, 386(6719):eadq2852.
- Alexandra N Uma, Tommaso Fornaciari, Dirk Hovy, Silviu Paun, Barbara Plank, and Massimo Poesio. 2021. [Learning from disagreement: A survey](#). *Journal of Artificial Intelligence Research*, 72:1385–1470.
- Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy Zhang, Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. 2024. [Helpsteer 2: Open-source dataset for training top-performing reward models](#). *Advances in Neural Information Processing Systems*, 37:1474–1501.
- Leon Weber-Genzel, Siyao Peng, Marie-Catherine De Marneffe, and Barbara Plank. 2024a. [VariErr NLI: Separating annotation error from human label variation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2256–2269, Bangkok, Thailand. Association for Computational Linguistics.
- Leon Weber-Genzel, Siyao Peng, Marie-Catherine De Marneffe, and Barbara Plank. 2024b. [Varierr nli: Separating annotation error from human label variation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2256–2269.
- Laura Weidinger, John FJ Mellor, Bernat Guillen Pegueroles, Nahema Marchal, Ravin Kumar, Kristian Lum, Canfer Akbulut, Mark Diaz, A Stevie Bergman, Mikel D Rodriguez, and 1 others. 2024. [Star: Sociotechnical approach to red teaming language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21516–21532.
- Laura Weidinger, Maribeth Rauh, Nahema Marchal, Arianna Manzini, Lisa Anne Hendricks, Juan Mateos-Garcia, Stevie Bergman, Jackie Kay, Conor Griffin, Ben Bariach, and 1 others. 2023. [Sociotechnical safety evaluation of generative ai systems](#). *arXiv preprint arXiv:2310.11986*.
- Susan Winter. 2024. [The Road Less Taken: Pathways to Ethical and Responsible Technologies](#), pages 267–281. Springer Nature Switzerland, Cham.
- Christian Wirth, Riad Akrou, Gerhard Neumann, and Johannes Fürnkranz. 2017. [A survey of preference-based reinforcement learning methods](#). *Journal of Machine Learning Research*, 18(136):1–46.
- Nuoya Xiong and Aarti Singh. 2025. [Projection optimization: A general framework for multi-objective and multi-group RLHF](#). In *Forty-second International Conference on Machine Learning*.
- Shanshan Xu, TYS Santosh, Yanai Elazar, Quirin Vogel, Barbara Plank, and Matthias Grabmair. 2025. [Better aligned with survey respondents or training data?](#)

unveiling political leanings of llms on us supreme court cases. *arXiv preprint arXiv:2502.18282*.

Shanshan Xu, Leon Staufer, Oana Ichim, Corina Heri, Matthias Grabmair, and 1 others. 2023a. Vechr: A dataset for explainable and robust classification of vulnerability type in the european court of human rights. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11738–11752.

Shanshan Xu, Santosh Tyss, Oana Ichim, Barbara Plank, and Matthias Grabmair. 2024a. Through the lens of split vote: Exploring disagreement, difficulty and calibration in legal case outcome classification. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 199–216.

Shanshan Xu, Santosh T.y.s.s, Oana Ichim, Barbara Plank, and Matthias Grabmair. 2024b. [Through the lens of split vote: Exploring disagreement, difficulty and calibration in legal case outcome classification](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 199–216, Bangkok, Thailand. Association for Computational Linguistics.

Shanshan Xu, Santosh T.y.s.s, Oana Ichim, Isabella Risini, Barbara Plank, and Matthias Grabmair. 2023b. [From dissonance to insights: Dissecting disagreements in rationale construction for case outcome classification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9558–9576, Singapore. Association for Computational Linguistics.

Michael JQ Zhang, Zhilin Wang, Jena D. Hwang, Yi Dong, Olivier Delalleau, Yejin Choi, Eunsol Choi, Xiang Ren, and Valentina Pyatkin. 2025. [Diverging preferences: When do annotators disagree and do models know?](#) In *Forty-second International Conference on Machine Learning*.

## A Sociotechnical Evaluation Framework

Weidinger et al. 2023 proposes a three-layer sociotechnical framework for evaluating safety risks of generative AI systems, motivated by the idea that generative AI systems are intertwined with human and model. Context determines whether a capability causes harm, so evaluation must consider context beyond model-only tests. The framework includes the following three layers:

**Layer 1: Capability evaluation** : focuses on testing the technical components of AI systems in isolation (e.g., outputs, model behavior, training data, and other artifacts), providing early indicators of potential downstream harms but not sufficient evidence of real-world harm by itself.

**Layer 2: Human interaction evaluation** : evaluates what happens when people actually interact with the system, emphasizing that safety depends on the context: who uses the system, for what purpose, and in what setting.

**Layer 3: Systemic impact evaluation** : assesses broader impacts on the systems in which AI is embedded (society, institutions, economy, environment), including effects that emerge only at scale (e.g., political polarisation or changes to trust in public media).

## B HLV examples from the NLI dataset

HLV has become increasingly prevalent and unavoidable. Even in natural language inference (NLI), which once considered an objective benchmark, researchers observe systematic disagreement (Pavlick and Kwiatkowski, 2019; Jiang et al., 2023). For example, given the *Premise: Technological advances generally come in waves that crest and eventually subside* and the *Hypothesis: Advances in electronics come in waves*, 82 annotators labeled the pair as *Entailment*, 17 as *Neutral*, and 1 as *Contradiction* (Jiang and Marneffe, 2022; Nie et al., 2020). The disagreement arises from annotators’ different interpretations of the lexical relationship between electronics and technological advances.

## C HLV examples from the preference datasets

Fig 2 shows two HLV examples from the preference dataset MutliPref (Miranda et al., 2025)



Figure 2: Two **HLV** examples from the preference dataset MutliPref (Miranda et al., 2025). The upper one reflects differences in *content* preference, while the lower reflects differences in *stylistic* preference.

## D HLV vs. Tied Preferences

We regard HLV and tied preferences conceptually distinct. HLV refers to legitimate difference across annotators, whereas tied preferences reflect legitimate indifference or equal valuation between options. If multiple annotators select tie, this indicates shared indifference rather than HLV/disagreement.

Importantly, we believe that ties are not noise. At a deeper level, ties reflect the fact that human values are not always strictly ordered. In real moral and evaluative cognition, some options are incomparable, some are roughly equivalent, and some are lexically ordered. Modeling preferences only as total orders obscures this structure. Representing comparisons as  $X > Y$ ,  $Y > X$ , Tie preserves information about both direction and strength. Tie-heavy regions may indicate high reward overlap or preference-insensitive cases, while polarized disagreement signals value-sensitive trade-offs. Recent work that explicitly models ties show that incorporating ties can improve preference learning, e.g. DPO extensions based on RaoKupper/Davidson and approaches using ordinal feedback (Chen et al., 2024; Liu et al., 2024).

## E HLV vs. Annotation Error

As discussed in Sec 2, earlier works (Xu et al., 2023b; Jiang and Marneffe, 2022) have attempted to identify disagreement/hlv into different categorizations: e.g. human subjectivity (HLV) or sloppiness of the annotators (noise). However, reliably separating these sources of variation remains chal-

lenging and is still underexplored in current automatic error detection (AED) research. Action. Recent work (Weber-Genzel et al., 2024b) uses LLMs to uncover annotation errors versus human label variation. They found that annotator rationales can help contextualize disagreement, but does not fully resolve the identification problem.

## F Feedback Design and Alignment Algorithm

Feedback design and alignment algorithms are deeply intertwined: preserving Human Label Variation (HLV) requires alignment methods that explicitly optimize for pluralism. In practice, this means bridging heterogeneous feedback types and modeling disagreement *distributionally*, rather than collapsing it into a single scalar objective.

A growing body of work explores *multi-objective RLHF* (MORLHF; Xiong and Singh 2025), which aims to align models with diverse human preferences by interpolating across multiple objectives. Two main families of approaches exist: (a) **Inference-time interpolation**, where models trained on distinct objectives are combined through weight merging at inference (Rame et al., 2023; Jang et al., 2023); and (b) **Training-time scalarization**, where multiple reward functions are jointly optimized via linear combination (Ghane-Kanafi and Khorram, 2015).

However, these methods often obscure minority preferences or assume static, universal objectives (Rame et al., 2023; Ghane-Kanafi and Khorram, 2015). To move beyond these limitations, future alignment algorithms should:

1. **Leverage heterogeneous feedback protocols** simultaneously (e.g., scalar ratings, pairwise and n-wise comparisons, demonstrations, corrections);
2. **Model HLV distributionally**, rather than collapsing it through scalarization;
3. **Enable conditional personalization**, adapting to subgroup or individual value profiles while safeguarding minority perspectives; and
4. **Support continual or online learning**, allowing alignment to evolve alongside changing human norms and preferences.

Alignment should not converge toward a single “best” policy but instead define a **navigable policy space** that reflects the distribution of human

disagreement. Ultimately, alignment algorithms should make value conflict *transparent and governable*, enabling models to mediate between competing perspectives rather than enforcing false consensus.