

Human-Centric NLP or AI-Centric Illusion?: A Critical Investigation

Piyapath T Spencer

Proceedings of the 38th Pacific Asia Conference on
Language, Information and Computation (PACLIC 38)

Nathaniel Oco, Shirley N. Dita, Ariane Macalinga Borlongan, Jong-Bok Kim (eds.)

2024

© 2024. Piyapath T Spencer. Human-Centric NLP or AI-Centric Illusion?: A Critical Investigation. In Nathaniel Oco, Shirley N. Dita, Ariane Macalinga Borlongan, Jong-Bok Kim (eds.), *Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation* (PACLIC 38), 773-779. Institute for the Study of Language and Information, Kyung Hee University. This work is licensed under the Creative Commons Attribution 4.0 International License.

Human-Centric NLP or AI-Centric Illusion?: A Critical Investigation

Piyapath T Spencer

Language and Information Technology Programme, Faculty of Arts, CU, Thailand
Center for Information and Language Processing (CIS), LMU Munich, Germany
linguistics@piyapath.uk

Abstract

Human-Centric NLP often claims to prioritise human needs and values, yet many implementations reveal an underlying AI-centric focus. Through an analysis of case studies in language modelling, behavioural testing, and multi-modal alignment, this study identifies a significant gap between the ideas of human-centricity and actual practices. Key issues include misalignment with human-centred design principles, the reduction of human factors to mere benchmarks, and insufficient consideration of real-world impacts. The discussion explores whether Human-Centric NLP embodies true human-centred design, emphasising the need for interdisciplinary collaboration and ethical considerations. The paper advocates for a redefinition of Human-Centric NLP, urging a broader focus on real-world utility and societal implications to ensure that language technologies genuinely serve and empower users.

only in the form of massive datasets used for training or in post-hoc attempts to align the model with human preferences. This approach, whilst impressive in results, arguably prioritises AI capabilities over addressing fundamental human communication needs or cognitive processes.

This paper argues that despite its name and stated intentions, much of what is labeled as Human-Centric NLP is, in fact, predominantly AI-centric. Rather than genuinely prioritising human needs and experiences, these approaches often incorporate human information primarily as a means to enhance AI performance. This misalignment between the proclaimed human-centric goals and the AI-centric reality has significant implications for the development, application, and societal impact of NLP technologies.

1 Introduction

“Human-Centric NLP” purportedly aims to develop language technologies that are more aligned to human needs, cognition, and behaviour (Hovy and Spruit, 2016). Some argue that by incorporating human factors into NLP systems, we can create more effective, ethical, and user-friendly language technologies (Jurgens et al., 2019; Yang, 2023; Kotnis et al., 2022; Wang et al., 2021). However, a critical examination of current practices and research trends in this field raises a provocative question: *Is Human-Centric NLP truly centred on human needs, or is it a mere AI-centric illusion?*

Consider, for instance, the development of LLMs like GPT-4 (OpenAI, 2024). Even though it appeared as a step towards more human-like language understanding through various tasks, these models primarily focus on improving performance metrics such as perplexity and accuracy on benchmark tasks. The human element often comes into play

Through a critical investigation of current research trends, methodologies, and case studies, the paper aims to address the AI-centric nature under the surface of human-centricity in NLP. The paper also examines how human data and behaviour are often exploited to improve NLP systems without necessarily addressing core human needs or concerns, as observed by Bender and Koller (2020) and Bender et al. (2021). Furthermore, The paper explores the ethical implications of this mischaracterisation and propose a framework for what “truly” human-centric NLP might entail, building on the work of scholars who have called for more genuine engagement with human factors in AI development (Crawford and Calo, 2016). Ultimately, this paper seeks to stimulate a re-evaluation of priorities in NLP research and development. It calls for a genuine shift towards human-centric approaches that place human needs, experiences, and well-being at the forefront, rather than treating them as mere tools for technological advancement.

2 The Promise of Human-Centric NLP

The concept of Human-Centric NLP emerged as researchers recognised the need to align language technologies more closely with human needs and cognitive processes, partly as a response to criticisms that they were too focused on technical performance metrics at the expense of real-world applicability and human factors. [Kotnis et al. \(2022\)](#) related NLP with the idea of Human-Centric Research (HCR) with the objective to “place *all* [human] stakeholders at the centre of research”. This paradigm shift promised to conduct research (and create technologies) that are more intuitive, ethical, and aligned with human cognitive processes and societal needs.

Ever since, advocates of Human-Centric NLP have made strong claims about its potential benefits. For instance, [Sap et al. \(2020\)](#) and [Kaushik \(2023\)](#) argued that incorporating human knowledge and reasoning patterns could lead to more robust and generalisable NLP systems. They suggested that such systems would be better equipped to handle the nuances and contextual complexities of human language. Moreover, Human-Centric NLP has been acted as a solution to ethical concerns in AI development. [Hovy and Yang \(2021\)](#) proposed that by centring human values and societal impact in the design process, we could create more responsible and fair language technologies.

The promise of Human-Centric NLP extends beyond improved performance. Researchers have argued that this approach could lead to more ethical and socially responsible AI systems. For example, [Bender et al. \(2021\)](#) in their influential paper on the dangers of large language models, emphasised the need for NLP research to centre on human values and societal impact.

As the field of Human-Centric NLP continues to evolve, researchers are exploring ways to balance technical advancements with ethical considerations and user-centred design. This approach represents a shift from purely performance-driven metrics to a more holistic view of NLP’s role in society. There is, nevertheless, ongoing work to translate these human-centric ideals into practical implementations across various NLP applications.

3 The Reality: AI-Centricity in Disguise

Whilst the concept of Human-Centric NLP promises an optimistic picture of language technologies aligned with human needs and values, a

closer examination of current practices reveals a different reality. Despite the rhetoric of human-centricity, many NLP systems and research directions continue to prioritise AI performance over genuine human-centric considerations.

The development and deployment of such large language models (LLMs) as GPT-series, from GPT-3 ([Brown et al., 2020](#)) to GPT-4 ([OpenAI, 2024](#)), serve as a prime example of this disconnect. These models have achieved impressive results on various NLP tasks, yet their development process and application raise serious questions about their alignment with human-centric principles. The data collection methods for these models often involve web scraping vast amounts of information without adequate consideration for privacy, consent, or representation issues. Once again, [Bender et al. \(2021\)](#) argue that this approach to data collection reflects a prioritisation of model performance over ethical considerations and diverse human perspectives.

Furthermore, the evaluation of these models primarily relies on performance metrics for benchmark tasks and leaderboards. As [Ethayarajh and Jurafsky \(2020\)](#) point out, these metrics often fail to capture real-world utility or alignment with human values, instead focusing on narrow technical capabilities. The emphasis on benchmark performance over real-world applicability as such shows the AI-centric nature of current NLP practices.

The resource allocation for training LLMs also reflects a focus on pushing the boundaries of AI capabilities rather than addressing specific human needs or environmental concerns. [Strubell et al. \(2019\)](#) highlight the immense computational resources required for training models, raising questions about the prioritisation of AI advancement over other important human-related considerations such as environmental impact or more targeted, human-centric applications of NLP technologies.

Beyond LLMs, the field of sentiment analysis provides another example of this disconnect. Tools developed for understanding human emotions often reduce complex affective states to simplistic binary (positive/negative) classifications to ease the computation instead of capture intricate human emotional experiences, making this a reductionist approach that reflects a preference for computational efficiency over truly capturing the complexity of human sentiment.

This gap between the stated goals of Human-Centric NLP and its practical implementation raises

critical questions about the field’s direction. Are we truly developing technologies that serve human needs, or are we simply creating more sophisticated AI systems that give the illusion of human-centricity? Critics like [Birhane \(2021\)](#) argue that the focus on technical advancements often overshadows crucial discussions about the societal implications of these technologies.

It becomes more clear that bridging this gap between the promise and reality of Human-Centric NLP requires a fundamental re-evaluation of priorities and practices in the field. The challenge lies in aligning the impressive technical capabilities of modern NLP systems with genuine human-centric principles that prioritise ethical considerations, user needs, and societal impact.

4 Some Cases to be Discussed

To further substantiate the critical examination of Human-Centric NLP, this section presents three cases that exemplify the correlations between human-centric pictures and AI-centric realities.

On Linguistic Varieties The first case study is [Ramponi \(2024\)](#), addressing the challenges of developing NLP technologies for the diverse language varieties in Italy. Although the author makes important arguments about the technological challenges of Italy’s linguistic diversity, his focus on NLP solutions overlooks crucial socioeconomic factors that influence language vitality and also pays insufficient attention to intergenerational transmission dynamics and the predominantly oral dialects, potentially marginalising these aspects in favour of written forms that are more amenable to current NLP techniques. Lacking of a clear framework for community-driven priorities raises questions about how speaker communities themselves might shape research agendas and tool development. Perhaps most tellingly, the approach, whilst interdisciplinary in intent, does not fully integrate insights from sociolinguistics, anthropology, and cultural studies – disciplines crucial for understanding the human dimensions of language use and preservation. Despite its aims, the study remains primarily anchored in a technology-first paradigm that may not fully capture or address the complex human realities of Italy’s linguistic landscape.

On Evaluation The second case study is [Ribeiro et al. \(2020\)](#). This paper introduces CheckList, a task-agnostic methodology for testing NLP models; whilst innovative in its approach to NLP model

evaluation, reveals several limitations in its human-centricity. The automated testing and model failures, as well as the predefined linguistic capabilities and test types, may not fully capture the nuanced, contextual nature of human language use and potentially oversimplify the complex, holistic nature of human communication. CheckList’s black-box testing approach which focusing on discrete linguistic phenomena risks perpetuating a disconnect between model development and the lived experiences of language users. The benchmark-centric view, contrasting differences between model performance and human-like understanding, doesn’t deeply explore how these issues relate to real-world language use. Furthermore, the user studies primarily focus on CheckList’s ability to generate more tests and uncover bugs, rather than on how it improves the user experience or addresses human-centric language needs.

On Human Data Our third case study is [Takmaz et al. \(2020\)](#), aiming to improve image captioning by incorporating human gaze data, ostensibly making the process more ‘human-centric’. The authors use eye-tracking data to guide the image captioning model, arguing that this approach better aligns with human attention patterns. Its heavy reliance on eye-tracking data as a proxy for human attention risks oversimplifying the complex cognitive processes involved in image interpretation and description. Although the study introduces sequential processing of gaze patterns, this approach also potentially oversimplifies the non-linear and iterative nature of human thought processes during image description tasks. Besides, the introduction of the SSD metric further demonstrates a focus on quantifiable outcomes rather than qualitative alignment with human linguistic behaviour. Notably, the paper gives limited consideration to individual differences such as cultural background, personal experiences, or emotional responses that significantly influence image interpretation. The paper apparently emphasises on improving AI performance through gaze data suggesting a prioritisation of technological advancement over a deeper understanding of human cognition. Apart from this, there features no discussion on real-world applications for this particular innovation, seemingly the AI-centric nature of the approach.

These case studies collectively demonstrate the ongoing challenges in achieving Human-Centric NLP. They suggest that *true* human-centricity re-

quires more than just improved performance metrics or the incorporation of human data. Instead, it demands a deep engagement with the complexities of human cognition, cultural contexts, and social dynamics.

5 Rethinking *Human-Centricity*

As critically examining the concept and implementation of Human-Centric NLP, several key questions emerge that need further discussion. These questions challenge the understanding of what it means for NLP to be truly human-centric and how it relates to broader concepts of human-centred design and real-world impact.

1. Is Human-Centric NLP Human-Centred Design?

Human-Centred Design (HCD) is an approach that puts human needs, capabilities, and behaviours at the forefront of the design process. Whilst Human-Centric NLP claims to prioritise human factors, it's debatable whether current practices truly align with HCD principles.

Traditional HCD involves extensive user research, iterative prototyping, and continuous user feedback (Harte et al., 2017). However, much of Human-Centric NLP research focuses on improving model performance on human-generated datasets or incorporating human-like features, rather than directly involving users in the design process. The case study on Italian language varieties (Ramponi, 2024) demonstrates this disconnect: whilst aiming to address human linguistic diversity, the approach remains largely technology-driven rather than user-driven.

To truly embody HCD, Human-Centric NLP might need to shift towards more participatory research methods, involving end-users throughout the development process, from problem definition to solution evaluation.

2. Does Human-Centric NLP use Human as Another Metrics/Benchmark?

There is a growing concern that Human-Centric NLP often reduces human factors to another set of metrics or benchmarks, rather than genuinely centring human needs and experiences. The Check-List methodology (Ribeiro et al., 2020) exemplifies this tension: it aims to test NLP models on human-like language tasks; however, it still fundamentally treats human language abilities as a one of the benchmark for AI performance.

Similarly, the study on gaze-guided image cap-

tioning (Takmaz et al., 2020) uses human eye-tracking data to improve AI performance, but it is questionable whether this truly captures the essence of human image interpretation or merely uses human behaviour as another optimisation target.

This trend risks oversimplifying the complexity of human language and cognition. A more genuinely human-centric approach might involve developing evaluation methods that go beyond performance metrics to assess the real-world utility and social impact of NLP systems.

3. Should Human-Centric NLP Take a Step Out from the Computer/Virtual World?

Human-Centric NLP often focuses on improving language technologies within digital environments. However, language is fundamentally a tool for human interaction in the physical world. There is indeed a pressing need for Human-Centric NLP to consider its impacts and applications beyond the virtual realm. Ramponi (2024) touches on this by addressing real-world linguistic diversity, but there is potential to go further notwithstanding.

This could involve studying the real-world consequences of NLP systems, such as their impact on human communication patterns or social dynamics. Developing NLP applications that bridge the digital and physical worlds, like improved assistive technologies for individuals with disabilities, considering the environmental and societal impacts of large-scale NLP models and infrastructures.

These discussions point upon the need for a fundamental re-evaluation of what constitutes Human-Centric NLP. Moving forward, the field should strive for a more holistic approach that truly embodies human-centred design principles, goes beyond using humans as mere benchmarks, and actively engages with the physical world implications of language technologies.

6 The Prospects for *Human Language Technologies*

To realise the potential of Human-Centric NLP, it is crucial to broaden its application across diverse domains. Whilst it is apparent that many current implementations even under the umbrella of HC-NLP focus on specific tasks or benchmarks, there is a significant opportunity to apply human-centric principles in such domains with wide-reaching societal impacts areas as healthcare, education, and

social justice. A prominent example for this sort of application is illustrated by [Antoniak et al. \(2024\)](#), outlining the necessity of ethical frameworks for utilising NLP tools within maternal healthcare as well as addressing critical issues such as clinician-patient power dynamics and systemic health disparities. The authors developed guiding principles focused on “contextual significance, holistic measurements, and valuing diverse voices” by directly engaging with affected communities. Aside from a methodological pattern for future research, they serve as an important resource for practitioners aiming to create inclusive and effective NLP technologies in maternal healthcare and beyond.

Inseparably linked with the *threat*, [Jonas \(1984\)](#) proposed that technology must be guided by a principle of responsibility, valuing long-term human welfare and ethical integrity. Human-Centric NLP faces inevitable ethical dilemmas, from biases in language models to the environmental impact of training large models. Addressing these challenges requires a commitment to transparency in model development, evaluation, and deployment. For NLP, this principle translates to a commitment to transparent trade-offs, recognising where and how models may fall short in meeting human-centred values and openly addressing the societal and environmental costs involved.

The path towards a truly Human-Centric NLP is undoubtedly challenging, especially, given resource limitations, the need for consistent model performance, and industry pressures for rapid deployment. It is, of course, understandable why a number of NLP researchers focus on achieving computational excellence. However, acknowledging these constraints does not preclude progress. Initial steps, such as introducing qualitative user-feedback mechanisms, incorporating human-centred metrics into model evaluations, or co-developing applications with end-users ([Lau et al., 2015](#); [Carreño and Winbladh, 2013](#); [Sreejith and Sinimole, 2024](#)), can make significant strides toward aligning NLP with human-centric ideals. Additionally, embedding ethical reviews in the development process, where models are evaluated for social implications before deployment, would set a standard of responsibility.

As NLP continues to evolve, its future trajectory depends on whether the field can transition from a narrowly AI-centric focus to a genuine commitment to human relevance, ethical responsibility,

and social accountability. The ultimate vision for Human-Centric NLP is not only to create models that excel at language processing but to foster tools that respect and enhance human agency, preserve cultural diversity, and engage responsibly with global issues. The question remains: *Will NLP’s legacy serve human welfare and societal well-being, or will it reinforce AI-centric illusions at the expense of human values?*

A Human-Centric NLP approach ought to position language technologies as extensions of human creativity, connection, and identity. The field can offer more than technological advancement through these humanistic principles; it can contribute to a world where language technologies genuinely serve as advocates to human expression, dignity, and community. This shift calls for a shared ethical commitment and the courage to prioritise human values over mere computational gains. In doing so, NLP can realise its potential to empower diverse voices, deepen understanding, and elevate the human experience in meaningful, lasting ways.

7 Conclusion

The paper have examined the disconnect between the promise of Human-Centric NLP and its current implementation. The analysis reveals that many so-called human-centric approaches in NLP remain fundamentally AI-centric nowadays. The case studies and subsequent discussion have demonstrated several key issues: the misalignment with true human-centred design principles, the reduction of human factors to mere benchmarks, and the limited consideration of real-world, physical impacts of NLP technologies. These findings indicate the need for a fundamental reframing of Human-Centric NLP. To address these challenges, we propose that truly Human-Centric NLP should embrace genuine human-centred design methodologies, develop holistic evaluation frameworks, expand its scope to consider broader societal implications, prioritise interdisciplinary collaboration, and centre ethical considerations throughout the development process. Only then can we hope to develop NLP systems that genuinely serve and empower humans in their diverse contexts and fulfil the true promise of human-centricity in NLP.

Acknowledgement

I would like to express my sincere gratitude to Professor Barbara Plank for her insightful comments

on this work and to the Research Affairs of the Faculty of Arts, Chulalongkorn University for providing generous financial support for this project.

References

- Maria Antoniak, Aakanksha Naik, Carla S. Alvarado, Lucy Lu Wang, and Irene Y. Chen. 2024. [Nlp for maternal healthcare: Perspectives and guiding principles in the age of llms](#). In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, page 1446–1463, New York, NY, USA. Association for Computing Machinery.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Abeba Birhane. 2021. [Algorithmic injustice: a relational ethics approach](#). *Patterns*, 2(2):100205.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Laura V. Galvis Carreño and Kristina Winbladh. 2013. [Analysis of user comments: An approach for software requirements evolution](#). In *2013 35th International Conference on Software Engineering (ICSE)*, pages 582–591.
- Kate Crawford and Ryan Calo. 2016. [There is a blind spot in AI research](#). *Nature*, 538(7625):311–313.
- Kawin Ethayarajh and Dan Jurafsky. 2020. [Utility is in the eye of the user: A critique of NLP leaderboards](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4846–4853, Online. Association for Computational Linguistics.
- Richard Harte, Liam Glynn, Alejandro Rodríguez-Molinero, Paul MA Baker, Thomas Scharf, Leo R. Quinlan, and Gearóid ÓLaighin. 2017. [A human-centered design methodology to enhance the usability, human factors, and user experience of connected health systems: A three-phase methodology](#). *JMIR Hum Factors*, 4(1):e8.
- Dirk Hovy and Shannon L. Spruit. 2016. [The social impact of natural language processing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.
- Dirk Hovy and Diyi Yang. 2021. [The importance of modeling social factors of language: Theory and practice](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 588–602, Online. Association for Computational Linguistics.
- Hans Jonas. 1984. *The Imperative of Responsibility: In Search of an Ethics for the Technological Age*. University of Chicago Press, Chicago.
- David Jurgens, Libby Hemphill, and Eshwar Chandrasekharan. 2019. [A just and comprehensive strategy for using NLP to address online abuse](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3658–3666, Florence, Italy. Association for Computational Linguistics.
- Divyansh Kaushik. 2023. [Robustifying NLP with Humans in the Loop](#). Ph.D. thesis, Language Technologies Institute, Carnegie Mellon University.
- Bhushan Kotnis, Kiril Gashteovski, Julia Gastinger, Giuseppe Serra, Francesco Alesiani, Timo Sztyler, Ammar Shaker, Na Gong, Carolin Lawrence, and Zhao Xu. 2022. [Human-centric research for nlp: Towards a definition and guiding questions](#).
- Rosa Lau, Fiona Stevenson, Bie Nio Ong, Krysia Dziedzic, Shaun Treweek, Sandra Eldridge, Hazel Everitt, Anne Kennedy, Nadeem Qureshi, Anne Rogers, et al. 2015. Achieving change in primary care—causes of the evidence to practice gap: systematic reviews of reviews. *Implementation Science*, 11:1–39.
- OpenAI. 2024. [Gpt-4 technical report](#).
- Alan Ramponi. 2024. [Language Varieties of Italy: Technology Challenges and Opportunities](#). *Transactions of the Association for Computational Linguistics*, 12:19–38.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.

- Maarten Sap, Vered Shwartz, Antoine Bosselut, Yejin Choi, and Dan Roth. 2020. [Commonsense reasoning for natural language processing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 27–33, Online. Association for Computational Linguistics.
- R. Sreejith and K.R. Sinimole. 2024. [User-centric evaluation of ehr software through nlp-driven investigation: Implications for product development and user experience](#). *Journal of Open Innovation: Technology, Market, and Complexity*, 10(1):100206.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and policy considerations for deep learning in NLP](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy. Association for Computational Linguistics.
- Ece Takmaz, Sandro Pezzelle, Lisa Beinborn, and Raquel Fernández. 2020. [Generating Image Descriptions via Sequential Cross-Modal Alignment Guided by Human Gaze](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4664–4677, Online. Association for Computational Linguistics.
- Zijie J. Wang, Dongjin Choi, Shenyu Xu, and Diyi Yang. 2021. [Putting humans in the natural language processing loop: A survey](#). In *Proceedings of the First Workshop on Bridging Human–Computer Interaction and Natural Language Processing*, pages 47–52, Online. Association for Computational Linguistics.
- Diyi Yang. 2023. [CS329X: Human Centered NLP](#).