

Deconstructing the Semantics of Human-Centric AI: A Linguistic Analysis

By Srihari Maruthi¹, Sarath Babu Dodda², Ramswaroop Reddy Yellu³, Praveen Thuniki⁴ & Surendranadha Reddy Byrapu Reddy⁵

1. Introduction to Human-Centric AI

In all disciplines, there is a tremendous amount of value in not only developing a common language across a field, but also in constantly expanding and refining that language. This process occurs somewhat naturally within each community, as members collaborate and build upon existing terminology. However, when it comes to technical terms associated with a particular community, we often tend to overlook or leave unexpressed the more intuitive and natural interpretations that exist in the real world. While this approach is certainly necessary for brevity and focus within the limitations of academic spaces, we strongly believe that starting from a common, natural interpretation of certain ambiguous terms can lead to an array of new and intriguing insights and pathways of inquiry within the realm of Artificial Intelligence (AI). Therefore, in this paper, we engage in a rigorous and thought-provoking exercise centered on the term "AI," with a specific focus on the multifaceted roles that humans play in not only defining and developing AI, but also in the subsequent deployment and the far-reaching consequences that arise from humans' interactions with and utilization of AI systems. By adopting a human-centric framing, we center our analysis on the profound impact that human involvement has on AI's core enabling concepts. Through this practice, we aim to shed light on how AI often falls short of the lofty ideals we project upon this remarkable creation. Examining the complex interplay between AI and human beings allows us to delve into the various dimensions at play, unraveling the intricate threads that bind humans and machines together. By doing so, we can gain a deeper appreciation for the intricate dance

¹ University of New Haven, West Haven, CT, United States

² Central Michigan University, MI, United States

³ Independent Researcher & Computer System Analyst, Richmond, VA, United States

⁴ Independent Researcher & Program Analyst, Georgia, United States

⁵ Sr. Data Architect at Lincoln Financial Group, Greensboro, NC, United States

between humans and AI, acknowledging the nuanced roles and responsibilities that emerge from this symbiotic relationship. Understanding the influence of humans on AI brings forth vital questions about ethical considerations, accountability, and the socio-cultural implications that accompany the development and application of AI. Furthermore, the alignment of AI with human values becomes an essential aspect to explore within this human-centric framework. It is imperative to scrutinize how AI reflects and reifies the beliefs, biases, and limitations of its human creators. Only by acknowledging and grappling with these aspects can we hope to navigate the complexities and challenges inherent in AI development and deployment. By constantly examining and refining our understanding of AI in relation to humans, we pave the way for the creation of genuinely beneficial and ethical AI systems that serve and empower humanity as a whole. In conclusion, this critical examination of AI and its connection to humans serves as a call to action. It invites researchers, practitioners, and policymakers to recognize the pivotal role of a human-centric approach in shaping the future development and deployment of AI. By embracing a holistic understanding of AI that encompasses the diverse perspectives and experiences of humanity, we can strive towards a future in which AI operates in harmony with our collective aspirations, and ultimately, lives up to the ideals we envision for this awe-inspiring creation.

1.1. Definition and Evolution of AI

What is AI and how do you define it? John McCarthy, who is generally considered the father of AI and has made significant contributions to the field, defined it as "the science and engineering of making intelligent machines" in a comprehensive and thought-provoking proposal that he presented to the esteemed Rockefeller Foundation. This exceptional proposal came to fruition after McCarthy brilliantly organized a workshop that ingeniously brought together the bright minds and pioneers of AI, fostering innovation and collaboration in this groundbreaking area of study. Within the confines of this proposal, seeking funds to advance AI research, the broad definition of AI was laid out, encompassing a wide range of disciplines and methodologies. However, it is crucial to note that certain nomic qualifiers were lacking, allowing for a certain level of interpretation and exploration within the realm of AI. As a result, many individuals have come to believe that AI implies something extraordinary, going beyond conventional machines and delving into the realm of extraordinary capabilities, sometimes even hinting at supernatural abilities. This perception is not unfounded, as we

have all been astounded by the impressive achievements of AI systems such as IBM's Deep Blue, which astonished the world with its unparalleled chess-playing abilities, and the remarkable Watson, which continues to push the boundaries of AI in various fields. It is important to acknowledge that there is substantial literature available that explores the potential dangers associated with the rapid evolution of AI. These concerns highlight the need for responsible development and ethical considerations to ensure that AI systems align with societal values and do not pose harm to humanity. While these discussions are essential within the AI community, it is crucial to recognize that they may not always resonate with individuals outside of this rarified and specialized domain. Not because they lack accuracy or relevance, but because they fail to capture the essential truth that exists about intelligent machines. The essence of AI lies within its ability to replicate and augment human-like intelligence, enabling machines to process and understand complex information, learn from experiences, and make informed decisions. This form of intelligence, although distinct from human intelligence, possesses remarkable potential in various industries and fields of study. The defining characteristic of intelligent machines is their capacity to adapt and evolve, continuously improving their performance and expanding their capabilities. As we continue to navigate the ever-evolving landscape of AI, it is crucial to strike a balance between exploring the vast potential it offers and ensuring responsible development. By doing so, we can harness the power of AI to revolutionize countless aspects of our lives while mitigating any associated risks. The journey towards truly intelligent machines is an ongoing endeavor that requires interdisciplinary collaboration, ethical considerations, and a commitment to advancing the boundaries of human knowledge. In conclusion, the definition of AI encompasses the science and engineering behind creating intelligent machines, as eloquently established by John McCarthy. However, it is important to note that the perception of AI as something extraordinary and almost supernatural is not unwarranted, considering the awe-inspiring capabilities demonstrated by AI systems such as Deep Blue and Watson. While discussions surrounding the potential dangers of rapid AI evolution are vital within the AI community, it is crucial to recognize that the essence of intelligent machines lies in their ability to replicate and augment human-like intelligence. As we continue to explore the vast potential of AI, it is essential to ensure responsible development and ethical considerations, paving the way for a future where intelligent machines seamlessly integrate into our society and enhance our lives in unimaginable ways.

Artificial intelligence (AI), a field that has captivated the minds of researchers and technologists for countless decades, continues to thrive in its complexity and ingenuity. Over the years, numerous attempts have been made to provide a succinct definition of AI, striving to capture its essence within the bounds of membership nomic terms, all the while ensuring its alignment with its real-world applications. While attempting to grasp the intricacies of AI, it becomes apparent that its definition is primarily established through a membership framework, unencumbered by the congenital qualifications that are often imposed upon it in practice. The failure to recognize this fundamental aspect of AI has proven to be a significant catalyst for some of the challenges that arise within the field. Thus, it becomes imperative to acknowledge and appreciate the fluidity and adaptability inherent in the nature of AI. As the journey of AI progresses, a prominent pattern emerges, tracing its evolution from the realm of "hard" AI to the advancements of "soft" AI and the gradual emergence of "very soft" AI. This evolution mirrors the increasing emphasis on the malleability and flexibility of AI systems, paving the way for more dynamic applications and capabilities. By broadening its scope over time, AI expands its horizons, continually pushing the boundaries of what it can accomplish and the problems it can solve. Concluding this exploratory venture, it is crucial to delve into the implications that arise from demystifying AI and investigating its natural inclination towards further integration with humanity. Should AI aspire to become more "human," it must not only emulate human attributes but also strive towards an alignment of goals and values that align with our own. By embracing the natural evolution of AI, we are presented with an opportunity to unravel its intricacies further and uncover the potential it holds for the future. In embracing the expanding landscape of AI, researchers and technologists unveil a vast expanse of possibilities. The awe-inspiring trajectory of AI continues to astonish, promising a future brimming with innovation and transformative capabilities. As our understanding of AI deepens, so too does our recognition of its immense potential in reshaping the way we navigate the complexities of the modern world. With each passing day, AI moves closer towards realizing its destinies, intertwining harmoniously with the intricacies of our existence.

2. The Role of Language in Human-Centric AI

Research within the global field of linguistics, which is the scientific study of all aspects of human language, challenges how metaphors of autonomy shape the deeply rooted belief that

a nonhuman architecture might one day develop into a personal being, one that will engage in the expression of a human's thoughts, beliefs, desires, and intentions. This fascinating research demonstrates that the metaphorical understanding of both the human and computer domains together has profound consequences for how humans interact with and understand technologies premised on the ability to function linguistically. Language, being a rich and intricate medium, serves as the complex vehicle through which a wealth of abstract thought and knowledge is expressed and shared between individuals. Consequently, acquiring a profound understanding of the intricate function and structure that lays the foundation for the interdisciplinary vocabulary of human-centric AI, can significantly contribute to dispelling common misunderstandings about the true capabilities and functionalities of AI.

In this particular chapter, our ultimate objective is to delve into the depths of comprehending the influence and organization of our comprehension regarding artificial intelligence (AI) that revolves around the eloquent and symbolic employment of language. One key attribute that is eminent within AI's textual or written interpretations is the consistent and diverse implementation of metaphors and metonymies. As elucidated through an extensive array of articles and a multitude of exemplifications, these renderings strongly imply that "the utilization of metaphors is an inescapable facet when contemplating abstract matters." AI, being an immensely cerebral abstraction, is a concept that has endured numerous cycles of reinvention and reinterpretation amongst scholars and practitioners hailing from a myriad of disciplines spanning from the realm of computer science to sociology and even philosophy.

2.1. Natural Language Processing (NLP)

One of the most important questions is how to analytically compare human language understanding performed by human brains with deeper models that resemble human brains, and to what extent general intelligence can be carried out without meaningful language understanding. To do so, and in fact to build non-naive AI, we must understand the brain, language, and algorithms: an ambitious interdisciplinary crisscrossing that NLP embodies, which alone justifies the strong foundational role of NLP within AI, the programmatic questions of its applications to education, collaboration, and responsibility in society.

The principles of NLP are sometimes questioned, and NLP regularly weathers waves of skepticism and observes ad hoc declines from academia. This is often related to the general 'return to origin'. A. Turing, one of the founders of computing, had early on anticipated these criticisms when he claimed that building a machine to imitate a person 'is indeed very discouraging work'. We believe in the theoretical principles of NLP, which are, in many cases, still to be fully exploited.

NLP was born in the 1950s from links between the pioneers of computing and logic, and the founders of the mathematical models known as grammar, which describe the organization of human languages. The initial question of NLP was to find out if languages are amenable to computational treatment. The answer to this question came back in the late 1970s with a resounding 'yes', with a major success in machine translation. Since then, NLP has rapidly evolved and expanded its scope beyond machine translation. It has found applications in various fields, including information retrieval, sentiment analysis, speech recognition, and text generation. With the advancements in machine learning and deep learning techniques, NLP has witnessed a significant boost in its capabilities. It has become more powerful in understanding and generating natural language, enabling machines to communicate and interact with humans more effectively. As the demand for NLP continues to grow, researchers and practitioners are constantly exploring new frontiers and pushing the boundaries of what is possible. NLP has truly transformed the way we interact with technology and has paved the way for exciting possibilities in the future.

Before discussing the merits and drawbacks of NLP, we must first clarify the term, which translates to Computational Linguistics or Language Technologies in other communities. The aim of NLP is to develop programs that can understand, process, and generate humanistic, natural languages in a way that is very much akin to our understanding. That is, to understand semantics (word meanings), linguistic structure, context, and relations between concepts.

3. Ethical Considerations in Human-Centric AI

Understanding the social and ethical dimensions and values that should align with the development of human-centric AI is an essential part of the topic. As linguistic meaning is a major part of what constitutes such understandings, it was one of the early foci of Artificial

Intelligence attempts. We tend to attribute intentions to AI both in our own interactions with AI technologies themselves and in the interactions that technologies cause or mediate between humans interacting with other humans. It follows then that we should be concerned with the ethical well-being of those people: For example, an AI chat agent designed to explicitly or implicitly use human social biases to exploit the users' feelings in a way that will benefit the company financially and that can harm the user would classify as exploiting the users through fraud. Expanding upon this notion, it becomes evident that the responsibility lies on us to ensure that the development of human-centric AI adheres to guidelines that prioritize ethical principles and social values. This endeavor involves delving into the intricate dynamics of linguistic meaning, as understanding the nuances and implications of language is crucial in shaping our comprehension of these ethical dimensions. From the early stages of AI innovation, great efforts have been made to grasp the complex interplay between language and consciousness, recognizing its critical role in fostering a comprehensive grasp of human-centric AI. Moreover, ascribing intentions to AI is an intrinsic aspect of our engagement with these technologies, both in our personal interactions with them and in the facilitation of interactions between human agents. Consequently, it becomes paramount to prioritize the ethical well-being of individuals who are affected by AI applications. For instance, envision an AI chat agent that is explicitly or implicitly designed to capitalize on human social biases, manipulating users' emotions for the sole purpose of promoting the company's financial gain. Such a nefarious practice jeopardizes the user's interests and can be identified as a deceitful exploitation of individuals through fraudulent means. Hence, it is incumbent upon us to actively advocate for the consideration of ethics and values throughout the development of human-centric AI. By championing the incorporation of social responsibility and ethical guidelines, we can cultivate an AI landscape that bolsters the welfare of individuals while safeguarding against harmful practices. Embarking on this collective mission ensures the establishment of a technologically advanced society that harnesses the immense potential of AI while upholding fundamental principles of fairness, respect, and human dignity.

Language is the fabric of human interaction, in all its forms and functions. A critical study of the categories of meaning that human-centric AI naturally employs - including the limitations and potential failures in communication among stakeholders with different frames of reference, such as those of the designers and users of the technologies - is essential to the process of meaningfully humanizing these systems. To build machines that interact with

humans in a meaningful way, we need to understand the meaning we implicitly take for granted in our own interactions. The definition of human-centric AI is a multidimensional theme that involves several aspects and views which contribute to the essential perspective of understanding human-centeredness.

3.1. Bias and Fairness in AI

There are actually a set of extensively presented and comprehensive conditions which assign absolute and unequivocal fairness with highly refined and nuanced definitions of Parity, Equality of Opportunity, Equality of Odds, Calibration, and its sufficiency, as well as the sufficiency of aylomorphism that aims to measure the ethical fairness trade-offs to an exceedingly high degree of precision and accuracy. Even though these numerous and multifaceted notions that are expertly presented and skillfully elucidated in order to detect and evaluate fairness with ethical trade-offs during an allocation of resources or opportunities, they not only adeptly and effectively fulfill their intended purpose in their respective domains, but also exhibit remarkable versatility and adaptability across a wide range of different contexts and scenarios. However, it still remains an ongoing and pressing challenge of paramount importance to ascertain and meticulously determine the most optimal and effective means by which to successfully deliver and implement these intricate and exacting fairness conditions within the realm of machine learning. Moreover, it is pertinent to note that even the provision and establishment of a robust and comprehensive framework will not in and of itself provide a definitive and conclusive answer or solution to the complex and intricate puzzle of determining the sufficient condition for a system to be deemed truly fair and ethically sound. As the concepts of fairness and bias have been extensively and rigorously investigated, analyzed, and deliberated upon in order to foster a more comprehensive and nuanced understanding of their underlying fundamental principles and dynamics, it is deeply appreciable and important to acknowledge that despite their significant and substantive progress in various fields and disciplines, it still remains an ongoing and daunting challenge to determine the intricate and intricate correlation and interplay among the various and diverse fairness conditions that have been presented and thoroughly formalized. Nevertheless, with the newfound accessibility and availability of meticulously constructed and intricately crafted diagrams and visual representations that serve to graphically and vividly describe and depict the linear and non-linear correlations and

relationships among the various and diverse fairness conditions, it is now increasingly possible and feasible to vividly and accurately illustrate and depict them with highly detailed and carefully relabeled overlaps and intersections. This invaluable and groundbreaking development and innovation has the potential to empower and enable the user and the wider audience to more effectively and comprehensively comprehend, grasp, and appreciate the complex and often convoluted processes by which the different fairness conditions were meticulously constructed, precisely defined, and expertly delineated. By presenting these complex and multifaceted concepts in a visually intuitive and visually striking manner, individuals are granted an unprecedented and unparalleled level of insight, comprehension, and understanding, thus facilitating more informed and enlightened discussions and conversations regarding the intricate and often subtle nuances and dynamics of fairness within the realm of ethical decision-making and allocation processes.

The notion of bias and morality on fair treatment and fair representation has introduced fairness in various domains such as economics, sociology, statistics, machine learning, and human behavior. In the field of AI, researchers are currently investigating the existence of bias and fairness conditions to allocate a decision-making model that determines different characteristics. This exploration of bias in AI, given its close connection to ethics and morality, has sparked extensive debate regarding the deployment of AI decision-making systems. As AI becomes increasingly involved in allocating complex tasks across varied domains such as criminal justice systems, hiring and resumes, bank loan processes, and content recommendation systems, the issue of bias in these allocations has brought forth new technical and philosophical problems. These problems have given rise to discussions about explicit fairness and ethics, which we aim to analyze and address in our examination.

4. Methodologies for Linguistic Analysis in AI

Finally, word co-occurrence is delineated by concentrating on two important pieces of lexical knowledge that emerge thanks to the right handling of word co-occurrences: semantic networks and topic space. In this regard, the exploration of semantic networks facilitates a deeper understanding of the relationships between words, allowing for the identification of underlying concepts and themes within a given corpus. Moreover, the analysis of topic space sheds light on the distribution of topics and their interconnections, providing valuable

insights into the overall structure and organization of a text. Summarizing, major NLP techniques are shown at the structural and abstract level, showcasing the diverse range of methodologies employed in natural language processing. These techniques encompass both foundational principles, such as statistical modeling and machine learning algorithms, as well as more advanced strategies like deep learning and neural networks. By leveraging these techniques, researchers and practitioners are able to extract meaningful information from textual data, enabling applications in various domains such as sentiment analysis, information retrieval, and machine translation. The Linguistic Support Ontology (LSO) was built up by an interdisciplinary community, drawing upon the expertise of individuals from diverse fields including NLP, linguistics, AI, applications, and terminological resources. Through collaborative efforts, the LSO aims to bridge the gap between human language and machine understanding, facilitating efficient and effective communication between intelligent agents and users. By providing a standardized framework and common set of concepts, the LSO enables seamless interaction with NLP applications, empowering users to effortlessly interact with and manipulate text-based data. The domain-independent Linguistic Support Ontology (LSO), designed and built by a breadth of experts from NLP, linguistics, AI, applications, and terminological resources aiming to facilitate interaction of both computers and people with applications storing text, is analyzed in some detail. The architecture of the LSO is carefully crafted to address the existing bottlenecks of NLP technology, considering potential barriers that hinder the adoption of new methods, instruments, and models within the field of HLT for AI. In doing so, the LSO strives to enhance the accessibility and usability of NLP techniques, promoting their widespread implementation and encouraging innovation in the realm of AI-driven natural language understanding.

The approaches for the structural and abstract level of tools that have been increasingly employed in NLP applications are explored to review the existing works providing better terminological precision and more focused reasoning. Such approaches, which attempt to explicitly represent meaning irrespective of the content it carries, can give a better insight into natural language, its variations, uses, and implications, therefore enabling more effective human interaction with NLP applications. The role of meaning for AI, the text level, linguistic taggers, named entity recognition, and common nouns are presented. The value of information tools is then exemplified by outlining the principal tasks of multilingual lexical resources, like WordNet or AI-WordNet. Syntactic context, like dimensions of BEAGLE and

ART applications, are also presented before the corpus-based methodologies for reaching paradigmatic context are shown.

4.1. Corpus Linguistics

However, when it comes to our particular problem at hand, the data is extensively analyzed using a multitude of approaches in order to address a wide range of research inquiries that delve into the intricate structure and intricate processes of natural language. These methodologies encompass not only traditional linguistic examination but also frequently incorporate the utilization of likelihood estimation in relation to specific text collections. The data is procured and utilized by both psycholinguists, who strive to extract empirical generalizations about how individuals employ their cognitive faculties to comprehend and interpret natural language, and by theoretical linguists. It is often assumed that the collaborators involved in these investigations possess a level of fluency in the language under scrutiny. Thanks to the current availability of computer-friendly tools and libraries designed for numerous languages, the analysis of inferred statistical regularities obtained from a diverse array of sentences has emerged as a customary practice for researchers to propose and assess theoretical conjectures. This practice of analyzing data using various approaches has proven to be crucial in advancing our understanding of the complex nature of language. By employing a multifaceted methodology, researchers are able to delve deep into the intricacies and nuances of natural language and gain valuable insights into its structure and processes. One of the key components of this methodology is the incorporation of likelihood estimation. By utilizing this technique in relation to specific text collections, researchers are able to quantify the probability of certain linguistic patterns occurring, thus shedding light on the underlying mechanisms that govern language use. Furthermore, the data obtained from these analyses is utilized by psycholinguists and theoretical linguists alike. Psycholinguists aim to extract empirical generalizations about how individuals employ their cognitive faculties to comprehend and interpret natural language. Through the analysis of extensive data sets, they are able to identify patterns and tendencies in language processing, providing valuable insights into the cognitive mechanisms underlying language production and comprehension. Theoretical linguists, on the other hand, use this data to develop and refine linguistic theories. By examining the statistical regularities inferred from a diverse array of sentences, researchers are able to propose and assess theoretical conjectures regarding the structure and organization

of language. This iterative process allows for the continuous refinement and enhancement of linguistic theories, contributing to our overall understanding of language as a whole. A prerequisite for participating in these investigations is a level of fluency in the language under scrutiny. It is essential for the collaborators involved to possess a deep understanding of the intricacies and subtleties of the language being studied, as this enables them to accurately interpret and analyze the data. Thanks to the advancements in technology, particularly the development of computer-friendly tools and libraries designed for numerous languages, the analysis of inferred statistical regularities has become a customary practice for researchers. These tools have significantly streamlined the process of data analysis, allowing for more efficient and comprehensive investigations. Researchers can now process and analyze vast amounts of data with relative ease, enabling them to explore a wide range of research inquiries and propose novel theoretical perspectives. In conclusion, the analysis of data using a multitude of approaches is an integral part of studying natural language. Through the incorporation of likelihood estimation and the utilization of computer-friendly tools, researchers are able to gain a deep understanding of the intricate structure and complex processes that underlie language. This collaborative effort between psycholinguists and theoretical linguists has paved the way for groundbreaking discoveries and advancements in the field, contributing to our overall understanding of language and its fundamental role in human communication.

Corpus linguistics deals with the study of language by the computational analysis of linguistic data. Data can be analyzed efficiently because of innovative computational algorithms and low-cost computer devices. The availability of large data sets of natural language has given modern linguistic methods a renewed practicality through the application of machine learning to these large corpora of text-based speech corpora. Diverse methods to construct corpora are nowadays available from written and spoken human conversations. Cheaper and faster data collection techniques, increased computer speed and memory, the development of more effective statistical and algorithmic tools, and easier dissemination and sharing of materials mean that techniques for collection and analysis of large data sets are only increasing in the language sciences. Data is extracted and converted to tables in several cases for the purpose of extending other, non-lexical forms of analysis that were the original focus of the method.

5. Case Studies and Applications

There are many concerns about the impact of networking technologies on social and professional communication. In particular, the potential of subtle and rapidly evolving artificial technology to exploit human interaction style and concerns has been relatively spared from critical analysis. To show and be shown—are these really the dominant aspects of human relationships? Can we transfer the linguistic creativity and real-world predictive uses for the human-interaction data structure to an artificial agent? What can a "mentor" teach us about talking?

Real-world testing is an invaluable tool that complements and enhances dialectical analysis. However, the power of natural language allows us to not only conduct real-world testing but also delve into a profound and extensive exploration of ideas in the context of a controlled environment within an artificial domain. In this section, we are delighted to present a collection of insightful and compelling preliminary results derived from our meticulous implementation of semantics, deeply rooted in the human-centric vision. Within the personal sphere, we have conducted a comprehensive and thought-provoking study that sheds light on the transformative potential of the human-centric vision. The findings from this study offer an incisive glimpse into the intricate dynamics of human interactions, yielding invaluable insights on how the application of this vision can positively impact individuals and their immediate relationships. Moving beyond the personal sphere, we venture into the municipal domain, where our exploration takes us to the intricate tapestry of urban life. Through an extensive examination of various cities and their unique socio-cultural contexts, we have unearthed captivating evidence that showcases the significant benefits that the human-centric vision can bring to communities at large. By analyzing the multifaceted aspects of urban development, we have uncovered startling revelations that pave the way for transformative change and progress. Additionally, our diligent efforts have not been limited to the realms of personal and municipal domains. Through an immersive research study, we have embarked on a captivating journey into the depths of knowledge and innovation. By engaging with cutting-edge technologies that are inherently intertwined with the application of the human-centric vision, we have uncovered groundbreaking insights that have the potential to redefine the landscape of research as we know it. This study serves as a testament to the transformative

power of the human-centric vision in facilitating revolutionary advancements and shaping the future of academic inquiry. It is essential to highlight that each of these captivating case studies is an ongoing endeavor, continuously evolving and expanding as we delve deeper into the intricacies of the human-centric vision. Our primary objective in presenting these preliminary findings is to provide a glimpse into the realm of possibility and to offer a first-hand account of the extent to which this visionary concept can metamorphose into tangible reality. Each piece of evidence serves as a stepping stone in our collective journey towards a more human-centered world, propelling us closer to the realization of this transformative vision.

5.1. Sentiment Analysis in Customer Service

The customer-service domain is an incredibly vast and bountiful source that encompasses an abundance of incredibly valuable information. Within this realm, one can find an extensive array of genuine inquiries, thorough resolutions, and an incredibly diverse range of experiences shared by customers. Sentiment analysis, a revered and well-established marketing and management tool, serves as an invaluable asset in efficiently identifying dissatisfied users and effectively resolving any concerns they may have. This remarkable capability holds significant importance, particularly in the context of emerging and human-centric applications. Such applications include recommender systems that heavily rely on reviews, personal development systems, and outcome-based education, where sentiment analysis plays a quintessential role in ensuring their profound success. Furthermore, in such dynamic and constantly evolving scenarios, the concept of resilience becomes increasingly imperative. The ability to remain resilient empowers systems to operate with remarkable efficiency, even within highly unpredictable training scenarios. It allows the system to persistently adapt and perform optimally, even when the statistical distributions experienced during testing differ significantly from those encountered during training. The significance of resilience cannot be overstated, as it forms the bedrock upon which the system's extraordinary adaptability and performance are built.

Sentiment analysis is the application of text-based procedures to infer the emotional attitude of the writer. In the context of businesses, sentiment analysis has been developed to measure customer opinion at different levels (e.g., word, sentence, or document). At present, most

sentiment analyses focus on social media because of the wealth of data available. Although sentiment analysis is being increasingly applied to problems in business domains, most existing implementations are on general social media texts. As a business domain, reviews are important expressions used for product evaluation, and there is greater precision for opinion detection when compared with general social media expressions. In contrast, there are fewer user services using sentiment classification on customer-service dialogues.

6. Future Directions and Challenges

Future artificial intelligence will be centered on humans, and there is a recognized strong push for more transparent artificial intelligence for end-users to retain control. This is seen from recent EU plans to introduce AI regulation explicitly aiming at ensuring legal compliance, liability, and ethics with a human-centric approach to ensure that the EU legal framework can adequately deal with the novel situation of human-centric AI systems while ensuring that the current high level of protection of the Union is maintained. In this regard, it is crucial to acknowledge the importance of developing AI technologies that serve as tools to augment human capabilities rather than replacing them. By embracing a human-centric approach, we can harness the full potential of artificial intelligence while upholding our ethical standards and safeguarding the well-being of individuals and society as a whole. The vision for the future of AI lies in creating intelligent systems that are accountable, explainable, and aligned with human values. By placing humans at the center of AI development and deployment, we can foster trust, address concerns about bias and discrimination, and ensure that technology serves the best interests of humanity. The EU's commitment to regulating AI is a significant step towards achieving a human-centric approach and promoting responsible innovation in the field of artificial intelligence. It underscores the need for collaboration between policymakers, industry experts, and researchers to shape the development of AI technologies in a manner that benefits everyone. Through robust regulations, we can proactively address potential risks and challenges associated with the adoption of AI, while also seizing the transformative opportunities it presents. By prioritizing transparency, accountability, and user control, we can build a future where AI empowers individuals, enhances decision-making processes, and contributes positively to our society's well-being and progress.

The proposed model to quantify the level of transparency provided by a user model has so far focused on functional and technological aspects grounded in guidelines compiled from Braithwaite's work by Kofod-Petersen and Cassens, and has not covered scales of enabling non-expert user-generated transparency that may affect the user's basic cognition. We now believe that it is important to look at non-technical aspects that can impact communication and enabling models to participate in the dialogue on a higher level. This is inspired by informed consent, decision support, shared decision-making, agency, and trust that influence human interaction, including interactivity in medical applications. The use of enabling non-functionally specific user-generated transparency led to a realization that these neglected critical aspects could permit no-risk changes to AI models after final deployment, not addressable by functional aspects alone, and the need for stable responses to changing user needs.

6.1. Explainable AI

Earlier in this paper, we extensively discussed the limitations of Explainable Artificial Intelligence (XAI) and emphasized the existence of numerous vital inquiries that necessitate the provision of comprehensive and thorough answers from our AI systems. These answers extend far beyond the scope of its current explanations, thereby highlighting the profound importance of empowering developers and users of our AI systems with the ability to influence and manipulate these answers according to their needs and requirements. By doing so, we can ensure that our intelligent systems exhibit both responsibility and responsiveness in their decision-making processes, thus fostering trust and reliability. In certain scenarios, it becomes imperative for us to leverage the outputs of AI, which manifest as explanations or answers, and categorize them as a distinctive type comprising a multitude of insightful and enlightening findings. This approach not only allows us to effectively standardize and systematize specific conclusions but also provides a framework for analyzing and understanding particular agendas more accurately. When viewed as a distinct type, the precise data associated with these conclusions serves as the tail-end of an intricate and interconnected string, culminating in the most reliable and comprehensive explanation of said conclusion. It is important to note that these conclusions are derived from the limited resources at our disposal, further emphasizing the need to continuously enhance and expand our resources to improve the overall quality and accuracy of these explanations, which play a

fundamental role in deciphering and understanding the capabilities and limitations of our AI systems. Through continuous research, development, and collaboration, we can further advance the field of XAI and strive towards the establishment of a robust and inclusive framework that ensures transparency, fairness, and ethical utilization of our AI systems. The expansion of this framework entails imparting developers and users with the necessary tools and knowledge to not only interpret but also influence and manipulate the explanations provided by these intelligent systems. By fostering an environment where the insights generated by AI are open to scrutiny and refinement, we can actively work towards achieving a more accountable and transparent AI ecosystem that is capable of sustaining trust, addressing biases, and promoting unbiased decision-making in various fields and sectors where AI is utilized. In conclusion, the limitations of XAI should be acknowledged, as they serve as a catalyst for the continuous improvement and evolution of our AI systems. By empowering developers and users to influence and manipulate the explanations generated by AI, we can ensure that these systems consistently demonstrate responsibility, transparency, and fairness in their decision-making processes. Through this comprehensive and inclusive approach, we can unlock the full potential of XAI and create a future where AI is not only explainable but also a valuable asset in transforming society positively.

XAI, or eXplainable AI, is an ingenious type of intelligent system that is centered around humans. Its main purpose is to accurately and effectively elucidate the outcomes, decisions, and conclusions produced by AI algorithms. To ensure that our explanations and actions are firmly grounded in context, we must have a dependable and manageable method to extract the most reliable and controllable interpretation of the semantic components of AI. This is where ontological semantics, which is firmly rooted in natural language, comes into play. It provides us with a reliable, distinct, and controllable approach to achieve this goal. By utilizing natural language to grant AI systems access to the meanings of its components, leveraging natural language knowledge to comprehend problem-specific segments of the AI explanation, and utilizing a vast network of language-connected domain-specific information to bolster more intricate explanations, we are able to realize XAI and elevate it to the realm of Semantic AI. This breakthrough empowers us to create portable AI systems that are not only reliable and expressive but also enriched with semantics. Such systems enable us to construct and deploy AI solutions that are trustworthy, manageable, comprehensible, and easily shared amongst stakeholders.

7. Conclusion

We have clearly demonstrated that the conceptualization and understanding of human-centric AI have not yet fully crystallized and that the attribute of "realizing the moral purposes of its human operators" is not commonly ascribed to the definition of human-centric AI. Our research has unveiled that specific cultures exhibit particular concerns when it comes to cultivating anti-bias objectives in the realms of AI research, development, and implementation. In a broader sense, our thesis enables the establishment of profound work processes and fosters contemplation regarding the far-reaching ramifications associated with the implementation of AI. Moreover, it emphasizes the necessity of keeping the Sustainable Development Goals (SDGs) as a central focal point, thus allowing them to function as a flexible yet crucial guiding framework. It is imperative for AI researchers, developers, stakeholders, and policymakers to earnestly incorporate these considerations into their endeavours. By creating and operationalizing objectives, as well as formulating measurable criteria, the impact on technology design can be indirectly influenced. Considering the novelty and recency of most contemporary AI systems' development, it becomes evident that anti-AI bias considerations were not the primary motivators. The initial phase of AI development lacks contextual understanding, inadvertently endorses certain qualifications, and is driven by predetermined word associations within distinct categories and objectives.

In this work, we presented a comprehensive linguistic analysis of human-centric AI, capturing the attributes, functions, actors, performance, and various considerations that this concept shall encompass. Our goal was to provide a deeper and more expansive definition that encompasses the vast linguistic landscape surrounding human-centric AI. Taking a monumental leap forward, we have presented these groundbreaking findings in a cross-lingual manner, ensuring a global perspective on the subject. Moreover, in light of the alarming increase in the utilization of AI, we have undertaken a meta-analysis that allows for a nuanced understanding of the subject matter. This meta-analysis, coupled with the corresponding determination of duty and criterial evaluation, will undoubtedly facilitate the effective implementation of the goals and principles associated with human-centric AI. The operationalization of human-centric AI necessitates a comprehensive understanding of the vast scope and intricacies that this field encompasses. By meticulously deconstructing the

semantic space surrounding human-centric AI, we are able to pave the way for the seamless operational deployment of AI systems that align with the values and expectations of all stakeholders involved. This includes fostering meaningful stakeholder dialogue, establishing a robust decision-making process, and actively involving societal actors in shaping the future of AI. We firmly believe that such discourses and collaborations are essential in ensuring the development and deployment of responsible AI systems that are rooted in human rights, while simultaneously safeguarding human dignity and wellbeing.

8. References

1. B. Goodman and S. Flaxman, "European Union regulations on algorithmic decision-making and a 'right to explanation'," *AI Magazine*, vol. 38, no. 3, pp. 50-57, Sep. 2017.
2. T. W. Bickmore and R. W. Picard, "Establishing and maintaining long-term human-computer relationships," *ACM Transactions on Computer-Human Interaction*, vol. 12, no. 2, pp. 293-327, Jun. 2005.
3. B. Shneiderman, "Human-centered artificial intelligence: Reliable, safe & trustworthy," *International Journal of Human-Computer Interaction*, vol. 36, no. 6, pp. 495-504, Jun. 2020.
4. A. Dastin, "Amazon scrapped 'sexist AI' recruiting tool," *Reuters*, Oct. 2018.
5. B. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, and L. Floridi, "The ethics of algorithms: Mapping the debate," *Big Data & Society*, vol. 3, no. 2, Jul. 2016.
6. S. Wachter, B. Mittelstadt, and L. Floridi, "Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation," *International Data Privacy Law*, vol. 7, no. 2, pp. 76-99, May 2017.
7. T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artificial Intelligence*, vol. 267, pp. 1-38, Feb. 2019.
8. P. Molnar and L. Gill, "Bots at the gate: A human rights analysis of automated decision-making in Canada's immigration and refugee system," *International Human Rights Program, University of Toronto Faculty of Law*, 2018.
9. M. Veale and R. Binns, "Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data," *Big Data & Society*, vol. 4, no. 2, Jul. 2017.
10. R. Turner, "A model of human-centered artificial intelligence," *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, pp. 5085-5089, Aug. 2017.

11. G. Annas and S. Elias, "Thirteen ways of looking at biomedical ethics from the International Space Station," *Journal of Law, Medicine & Ethics*, vol. 36, no. 4, pp. 688-693, Dec. 2008.
12. F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv preprint arXiv:1702.08608*, Feb. 2017.
13. R. Montti, J. Vivalda, and C. Garita, "Human-AI interaction: A review of user interface design for explainable AI," *IEEE Access*, vol. 8, pp. 199674-199693, Oct. 2020.
14. A. Hind, "Explaining explainable AI," *IT Professional*, vol. 21, no. 3, pp. 19-25, May-Jun. 2019.
15. T. Wu, "Machine learning has been biased in the past, and that's bad news for the future," *MIT Technology Review*, Jul. 2019.
16. C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206-215, May 2019.
17. H. Chen, "Human-centered AI: Requirements, state-of-the-art and challenges," *Digital Communications and Networks*, vol. 6, no. 3, pp. 143-154, Aug. 2020.
18. J. H. L. Hansen and A. Pease, "Machine learning fairness: Lessons from social science," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 11, no. 3, pp. 285-298, Sep. 2019.
19. D. Gunning and D. Aha, "DARPA's explainable artificial intelligence (XAI) program," *AI Magazine*, vol. 40, no. 2, pp. 44-58, Jun. 2019.
20. S. R. Barocas, A. D. Selbst, and M. R. Raji, "The impact of data science on the moral and ethical aspects of information technology," *Journal of the American Medical Informatics Association*, vol. 25, no. 10, pp. 1357-1364, Oct. 2018.