

Socially Aware Language Technologies: Perspectives and Practices

Anonymous ACL submission

Abstract

Language technologies have made enormous progress, especially with the introduction of large language models (LLMs). These advances can, however, exacerbate a variety of issues that models have traditionally struggled with, such as bias, evaluation, and risks. In this perspective paper, we argue that many of these issues share a common core: a lack of awareness of the factors, context, and implications of the social environment in which NLP operates, which we call **social awareness**. While NLP is getting better at solving the linguistic aspects, relatively limited progress has been made in adding the social awareness required for language applications to work in all situations for all users. Integrating social awareness into NLP models will make applications more natural, helpful, and safe, and will open up new possibilities. Thus we argue that substantial challenges remain for NLP to develop social awareness and that we are just at the beginning of a new era for the field.

1 Introduction

Natural language processing (NLP) has made significant strides in recent years, thanks in part to the introduction of large pretrained language models (LLMs) based on Transformers (Vaswani et al., 2017; Brown et al., 2020). As a result, performance on various NLP tasks, such as machine translation, sentiment analysis, or conversational agents, has significantly improved. These models now perform these tasks seemingly as well as, if not better than, humans (Tedeschi et al., 2023). On the other hand, a growing number of issues and shortcomings with these models has been reported. Some of these issues include bias (Bolukbasi et al., 2016), toxicity (Gehman et al., 2020), trust (Litschko et al., 2023), and concerns about fairness (Hovy and Spruit, 2016; Blodgett et al., 2020; Shah et al., 2020; ElSherief et al., 2021). Word embeddings, which represent words in a mathematical space,

can, for example, inadvertently capture and reinforce biases in training data, perpetuating stereotypes and inequalities (Bolukbasi et al., 2016; Gonen and Goldberg, 2019). Machine translation systems have been shown to generate translations with unintended biases or inaccuracies (Vanmassenhove et al., 2018; Hovy et al., 2020), potentially exacerbating cultural and societal misunderstandings (Bird and Yibarbuk, 2024). Furthermore, NLP applications are still insufficient for tasks that require social awareness, especially in high-stakes areas like health care. In its current form, NLP only serves a subset of people and situations that use language technology (Held et al., 2023).

Many of these issues facing modern NLP share a common core. Namely, they result from failing to consider language (technologies) in the context of communities, cultural and ideological differences, and social contexts. All these social awareness aspects are relevant not just for English, but also the 7,000 languages out there (Joshi et al., 2020), adding complexity to the problem. These issues fall under social awareness:

Social Awareness refers to the ability to be aware of social factors, contexts, and social dynamics, as well as the implications of language use on the broader social environment.

Social awareness is undervalued in current NLP. Very traditional NLP models often focus on syntax, grammar, and lexicon, but have not made much progress in capturing cultural context and social interactions. The inherent difficulty of operationalizing and integrating these complexities into today’s LLMs is a significant reason. We argue that we need to address this issue to take NLP to the next level. We must broaden the scope of NLP technologies to a wider range of people and situations, and advance and promote inclusivity and accessibility across different languages and cultures (Hovy and

082 Yang, 2021; Hershcovich et al., 2022). This is a
083 message that needs to be **repeated** in the NLP com-
084 munity because we are nearing a point where such
085 social awareness can be integrated into systems.

086 Social awareness is not restricted to NLP; it
087 should be an integral and foundational component
088 across all modalities of AI. While our view ap-
089 plies most readily to NLP, it is also relevant for
090 vision (Fathi et al., 2012) and robotics (Breazeal,
091 2003), for example. Social awareness governs the
092 dynamics of human-human and human-AI interac-
093 tions and has an impact on knowledge acquisition
094 and use. Language, as a means of communication,
095 serves as a tool for individuals to achieve a variety
096 of goals, even though it is generated and consumed
097 by people from various backgrounds. NLP’s po-
098 tential insights and applications will inevitably be
099 limited if it does not consider the interaction of in-
100 dividuals, the context in which language is uttered,
101 and the specific goals it should achieve. Knowledge
102 of such goals or capabilities in turn enables users to
103 gain more trust in NLP system— social awareness
104 is also an important factor for more trustworthy
105 NLP in the future (Litschko et al., 2023). This
106 is because language is more than just words and
107 grammar; human society and culture is inextricably
108 linked to it. By modeling the social factors that
109 influence language, our AI systems can better un-
110 derstand and connect with people, and expand their
111 scope and depth. Concretely, in this position paper,
112 we introduce three key aspects socially aware NLP
113 needs to account for to work, namely **social factors**
114 (Section 2.1), **social interaction** (Section 2.2), and
115 **social implication** (Section 2.3).

116 2 What Is Socially Aware NLP?

117 Pentland (2005) defines *socially aware computa-*
118 *tion* as systems that can understand social signal-
119 ing and social context. The author argues that fo-
120 cusing on such dimensions can improve collective
121 decision-making and keep users informed. From a
122 psychology perspective, Daniel Goleman defines
123 a related term—*emotional intelligence*,—which
124 he breaks down into four subsets: self-awareness,
125 self-management, social awareness, and relation-
126 ship management (Hernez-Broome, 2012). For
127 emotional intelligence, social awareness requires
128 the ability to accurately understand other people’s
129 emotions and empathize with them, which relates
130 to having a Theory of Mind of others (Tomasello,
131 2014; Premack and Woodruff, 1978).

132 The concept of **social awareness** in natural lan-

guage understanding refers to shifting focus away
133 from classic tasks and benchmarks and instead en-
134 couraging tasks, models, and evaluations to con-
135 sider social factors (Hovy and Yang, 2021), so-
136 cial context, and social dynamics communicated
137 through language. The “awareness” pertains to
138 language technologies themselves being designed
139 to recognize and exhibit such social aspects and
140 process these socially driven meanings and impli-
141 cations behind language as humans do. In addition
142 to the increased awareness of language technolo-
143 gies, researchers or practitioners should also be
144 aware of these social aspects to design language
145 technologies that are socially aware. 146

147 **Socially aware language technologies**

148 refer to the study and development of
149 language technologies from a social
150 perspective, to provide NLP systems
151 with the ability to understand the social
152 context, perspectives, and emotions ex-
153 pressed in language by humans.

154 In other words, a socially aware system should
155 demonstrate emotional intelligence, social intelli-
156 gence, cultural competence, or perspective-taking
157 abilities. Consequently, research in socially aware
158 language technologies will focus on developing
159 new algorithms, models, evaluation metrics, and
160 approaches that allow NLP systems to be more so-
161 cially aware and to better recognize and respond
162 to social cues, cultural nuances, and other human
163 communication-related factors. This research also
164 entails the design and implementation of socially-
165 aware NLP systems for use in real-world applica-
166 tions such that their use, implications, and impact
167 are all understood during their development.

168 2.1 Social Factors

169 Hovy and Yang (2021) lay out seven *social fac-*
170 *tors* that NLP systems need to incorporate: *speaker*
171 *characteristics*, *receiver characteristics*, *social rela-*
172 *tions*, *context*, *social norms*, *culture and ideology*,
173 and *communicative goals*. These aspects provide
174 an easy-to-use conceptual taxonomy on different
175 social factors in language. It is also essential to con-
176 sider the incorporation and understanding of social
177 factors in various domains, including psychology,
178 sociology, and more.

179 **Social factors** refers to a wide range of
180 social aspects that shape the way we un-
181 derstand language use, including but not

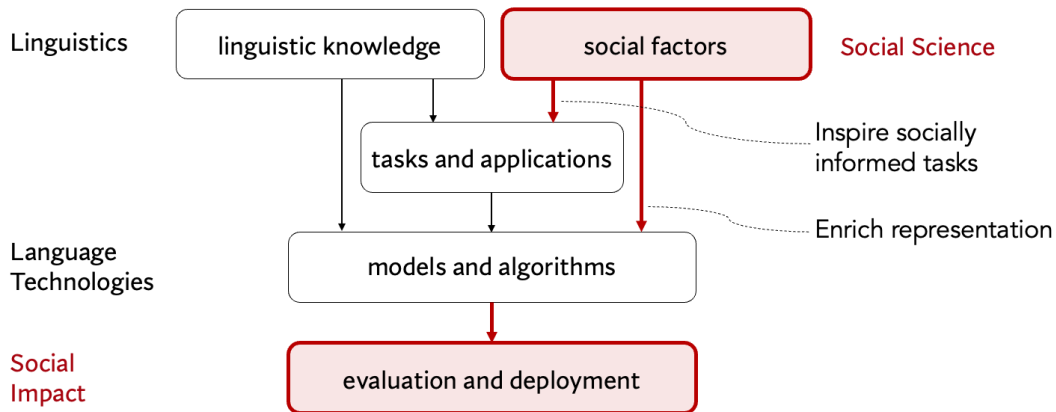


Figure 1: How social factors relates to language technologies, linguistics and social sciences

limited to who is the speaker; who is the receiver; what is their social relation; in what context; guided by what kinds of social norms, culture, and ideology; and for which communicative goals.

Figure 1 (without the highlighted red boxes) shows how current NLP works. We typically have some tasks inspired by linguistic knowledge and then build and evaluate models for the task, such as Natural Language Inference (Bowman et al., 2015). To develop socially aware language technologies, we have to incorporate social factors from social science into this pipeline. Social factors, in particular, can inspire socially informed tasks through additional objective functions and tasks. Operationalizing social phenomena will augment the current pool of tasks to better reflect users’ needs and likely lead to increased user trust. Furthermore, social knowledge can supplement existing representations in models (Nguyen et al., 2021). Social signals may provide alternative supervision for representation learning and next-word prediction. Current models have internal representations of social factors but do not appear to actively draw on them (Lauscher et al., 2022). With social awareness integrated into the pipeline, the outcome can produce social impact, not only about the typical task evaluation metric but also the impact on people.

2.2 Social Interaction

To develop socially aware language technologies, we must go beyond social factors (Hovy and Yang, 2021) and consider a wide range of **social interactions** and their **social implications**, such as the relationship, organizational, and cultural norms that

govern interpersonal communication. While social factors intrinsically encompass language, culture, and behavior, the awareness of social interactions and implications digs into the *dynamic context* in which NLP systems operate and their impact.

In social science, many theories, such as those surrounding social influence and norms, define essential dimensions of social interaction, shedding light on the intricate processes underlying human behavior and interaction. The desire to communicate within this environment helps drive language development, and our environment also profoundly influences the speed and efficiency of our language acquisition (Plunkett, 1997). This perspective posits that language is not an isolated construct but emerges as a product of social exchange and communication, aligning closely with the interactionism paradigm in sociology (Snyder and Ickes, 1985). Social norms govern social behavior and are regarded as groups’ shared standards of acceptable behavior. Some social norms become laws and rules, while others remain informal but equally influential. Integrating social norms into language introduces a layer of complexity beyond mere vocabulary and grammar. The work of Lapinski and Rimal (2005) highlights the nuanced interplay between social norms and language, demonstrating that linguistic expressions often serve as vehicles for the expression and reinforcement of these norms. All of these provide a rich and multifaceted foundation for our explorations of the complicated space of social interaction.

Social interaction refers to the interaction dynamics, including social exchange between individuals, other people in the

context who are involved and other activities surrounding the context.

Similar to *how we see ourselves comes from our perception of how others see us*, language use is influenced by others' views of the language, especially as people increasingly interact with LLMs. Socio-technical NLP systems exist within an ecosystem of social interactions, where users, developers, and stakeholders come together to create, deploy, and use these technologies. Many factors influence these social interactions, including power dynamics (Prabhakaran et al., 2013), trust (Litschko et al., 2023), and user expectations (Dhuliawala et al., 2023). The design of NLP systems must consider how these social interactions shape user experiences (Jakesch et al., 2023; Liu et al., 2022b) and how they impact the adoption and effectiveness of the technology.

2.3 Social Implication

Social implication encompasses a complex interplay of positive and negative effects of NLP systems on society. Understanding the effects and consequences of language technologies on society is crucial to the responsible development and sustainable use of NLP, which includes but not limited to assessing biases and stereotypes (Dev et al., 2022); considering how systems affect global populations, not just those in the North (Song et al., 2023; Ranathunga and de Silva, 2022); investigating the dissemination of misinformation and dual use of NLP systems; examining concerns around job displacement (Eloundou et al., 2023) and understanding how model uses are associated with creativity, trust, and productivity (Chakrabarty et al., 2024; Liu et al., 2022b).

Social implication refers to broadly understanding an NLP system's implications and impact on society.

Aside from understanding such broad implications of NLP systems, social implications can also inspire model design, such as designing models to consider the implications of their outputs. For example, work in prompt safety can be seen as an initial step at imbuing models with a sense of what responses have harmful social implications and teaching them to decline to answer (e.g., Bianchi et al., 2023). As a society, we need to understand these social implications if we are to use NLP to its full potential and mitigate harms.




Figure 2: Conceptual structure of socially aware NLP: social factors, social interaction, and social implication. Note that this is **not an exclusive partition**, but one way for us to understand the scope of different tasks.

Figure 2 shows that in development, social awareness is first present in the initial design of the task and algorithms (social factors). It plays a role in the middle ground of interactions and activities around it (social interactions) and the outer layer of impact (social implication). By putting a strong emphasis on social factors, social interaction, and their social implications, we hope socially aware language technologies can foster development that improves communication and aligns with ethical and societal values, promoting a more harmonious and equitable sociotechnical ecosystem.

2.4 Emerging Work in Socially Aware NLP

An increasing amount of research has been conducted around socially aware NLP.

From a **social factor** perspective, research has attempted to understand speaker and receiver characteristics (Flek, 2020; Hovy et al., 2020), such as personalized text generation (Wu et al., 2021) or demographic aspects (Hovy, 2015); model social relations to improve language understanding (Yang et al., 2019b; Iyyer et al., 2016); analyze culture for more contextualized algorithms (Huang and Yang, 2023; Kozlowski et al., 2018; Garimella et al., 2019); align NLP systems with values and preferences (Durmus et al., 2023; Liu et al., 2022a; Bai et al., 2022); and investigate the broader goals behind language interaction (Stab and Gurevych, 2014; Yang et al., 2019a).

From a **social interaction** perspective, recent work around LLMs has examined how humans interact with LLMs to accomplish tasks (Yuan et al., 2022; Shaikh et al., 2023a), and to simulate human interactions (Liu et al., 2023; Park et al., 2023; Aher

et al., 2023), building on earlier studies that used how humans interact to develop social or linguistic theories such as common ground (Shaikh et al., 2023b; Li et al., 2023; Paranjape and Manning, 2021; Pilán et al., 2023).

From a **social implication** perspective, a growing amount of research has started to look at the implications of language technologies, such as how they affect the labor market (Eloundou et al., 2023; Chen et al., 2023), how they assist in the preservation of languages and dialects (Sap et al., 2019; Ziems et al., 2022b) (see Meta’s efforts to develop high-quality machine translation capabilities for most of the world’s languages)¹, and how they can be used to support education (Kasneci et al., 2023), crisis management (Goecks and Waytowich, 2023) and accessibility (Gadiraju et al., 2023).

3 What Is Not Socially Aware NLP

As with many other things, we can best identify social awareness by its absence. Without social awareness, NLP technology will disregard social or cultural taboos, fail to consider personalized aspects of language applications, use language that the target audience cannot understand (due to age, education level, or other factors), or respond with inappropriate or hurtful responses (e.g., telling a suicidal user to kill themselves (Dinan et al., 2022)). Here, socially aware NLP differs from personalization as it aims to incorporate a broader context of language use, such as broader social and cultural groups, while personalization focuses more on the individual for customized user experience.

Socially aware NLP and “*NLP in a social context*” are related but *not* the same concept. The use of NLP techniques to analyze and understand language use *in social settings* such as online communities, political discourse, and public opinion is referred to as NLP in a social context or sometimes through the text lens of Computational Social Science (CSS), which often develops NLP models in service of uncovering patterns and trends in text in order to answer questions in the social sciences. In contrast, *socially aware NLP* refers to NLP models’ ability to recognize and respond to social factors, context, and cultural nuances in human language *in all situations*. It entails developing NLP models that recognize and respond to social cues such as tone of voice, sarcasm, humor, and cultural references to provide more accurate and effective mod-

¹<https://ai.meta.com/research/no-language-left-behind/>

els for understanding and communicating.

Compared to NLP in social context, CSS has a broader scope in using computational approaches (e.g., via networks, text, simulation or experiments) to analyze and model social phenomena, as well as understanding human behaviors and social sciences (Lazer et al., 2009). The focus of CSS is primarily on modeling and understanding social phenomena through computation, whereas socially aware NLP aims to refine language technologies that take social and cultural considerations into account. The other related concept is theory of mind (ToM; Grant et al., 2017; Le et al., 2019; Sap et al., 2022), which refers to the ability for models to reason about the mental state (e.g., intents, emotions, or beliefs) of others. While ToM and socially aware NLP share an overarching goal of improving models’ ability to interact with humans in a more understandable and socially appropriate manner, ToM differs from socially aware NLP by attempting to mimic human-like understandings of others’ mental states by attributing mental states and understanding intentions.

Human-Centered NLP and socially aware NLP both emphasize how to make NLP aware of human factors and aligned with real-world needs and contexts, as well as concerned with ethical considerations such as fairness, privacy and equality. In contrast, human-centered NLP focuses more on user-centered design to create systems tailored to user needs and is often based on iterative design, usability testing, and human in loop approaches towards improved human-system interaction, while socially aware NLP focuses on making NLP more aware of social context in which language is used and social dynamics of communication such as bias, cultural norms, and societal impact.

Similar comparisons also apply to the difference between *human-like* and *social awareness*. Human-like aims to create the type of user experience that mimics human-to-human interactions by making human-AI interactions natural and familiar to humans, while social awareness *further* encourages such interactions to be appropriate and considerate of the social environment.

4 How to Build Socially Aware NLP

We have outlined that socially aware language technologies need to understand 1) social factors, 2) social interaction, and 3) social implication to function well. However, it is not trivial to incorporate these different aspects into the development pro-

cess of socially aware NLP.

4.1 Considerations for Socially Aware NLP

Building socially aware NLP requires a combination of technical and ethical considerations: **(C1) Access to diverse communities:** Socially aware language technologies need access to large and diverse datasets that reflect the linguistic and cultural diversity of the target user groups, which ensures that the models recognize and respond to social cue variations specific to different communities (Yin et al., 2021; Sharma et al., 2023; Wang et al., 2024). Consider standard languages or mainstream user groups, versus low-resource languages and dialects, and vulnerable populations such as older adults or people with cognitive impairments. **(C2) Incorporation of context and interaction dynamics:** Developing language technologies that are socially aware must incorporate context as well as interaction dynamics, including tone of voice, relevant contexts, and domain users involved in the interaction, rather than as static, standalone multiple-choice questions proxies for language understanding. **(C3) Ethical and social considerations:** Socially aware language technologies must be designed with ethical and social considerations, such as fairness, transparency, and privacy, to avoid perpetuating stereotypes or biases (Ma et al., 2023), and to respect user privacy. **(C4) Iterative design and continuous learning:** Socially aware NLP models must be continually monitored and improved to ensure they are effective and up-to-date with changing norms. This process may involve iteratively incorporating feedback from domain users and updating the models as needed to improve their accuracy and effectiveness.

4.2 Process of Building Socially Aware NLP

Let’s take it a step further and explore what it means to *build socially aware language technologies*. One can start with models that identify social factors such as the speaker, the receiver, and other social contexts within a given interaction. Many existing studies have take this path, such as personalization (Wu et al., 2021) and the inference of social relations (Iyyer et al., 2016). Big data or deep learning models *alone* might not be sufficient to enable socially aware NLP, as social awareness also requires careful consideration of interaction dynamics, ethical and social issues, as well as the potential impact of NLP models on society. Socially aware NLP should be able to handle a broader range of inputs than current models, and distinguish and interpret complex social nuances effectively. If these models

are successful in capturing various social implications, one could intuitively expect them to enhance downstream applications.

Socially aware language technologies may be required to not only process but also articulate the reasoning behind their predictions in a reliable manner (Sorensen et al., 2024). This requirement involves a layer of self-explanation, which explains the models’ internal workings and their assumptions about social factors as they make them (Weber et al., 2023; Rajani et al., 2019). Even more than traditional models, socially aware models require transparency to ensure fair and justified outputs. Current LLMs lack such properties, and their reasoning is often non-reliable (Turpin et al., 2023; Mondorf and Plank, 2024). Thus, it is generally best to design a system to be socially aware from the beginning (e.g. Kotnis et al., 2022) instead of adding it as a patch later on. However, it is still possible to update an existing system so that it is more socially aware if care is taken to identify improvement areas, incorporate social data and context, and test and refine the system based on user feedback.

4.3 Key Directions for Socially Aware NLP

To advance socially aware NLP, we need to measure its progress. NLP has established a rich culture of tasks. We introduce multiple example broad directions around socially aware NLP to foster new approaches (see Table 1): (1) *Formulating tasks that operationalize social awareness*, which could benefit from theories and insights from social science and other related fields as discussed in Section 2. Many tasks in NLP have started to model social awareness, such as formulating the task of identifying hate speech (ElSherief et al., 2021; Breitfeller et al., 2019) or recognizing social relations (Iyyer et al., 2016; Choi et al., 2021), etc. These tasks provide representations that can be integrated to improve socially aware language understanding. (2) *Developing computational methods that detect social awareness* The goal is to develop computational models that can detect social awareness signals, as there is often an inadequate amount of supervised data for social factors, and generic NLP methods might not work well. (3) *Building systems that exhibit social awareness* Socially aware NLP aims to produce social impact by integrating social awareness into the systems’ development process to consider diverse social factors. (4) *Evaluating social awareness* To evaluate interventions in the wild, we must consider in addition to quantitative

Task Description	Social Factors	Social Interaction	Social Implication	Considerations
(1) Formulate tasks that operationalize social awareness	✓			C1, C3
(2) Develop computational methods that detect social awareness	✓	✓		C2, C4
(3) Build systems that exhibit social awareness.	✓	✓	✓	C2, C3, C4
(4) Evaluate social awareness with real-world application	✓	✓	✓	C1, C2, C3
(5) Build socially aware NLP for real-world uses	✓	✓	✓	C1, C2, C3, C4
(6) Understand how socially aware NLP affects people and society	✓	✓	✓	C1, C3, C4

Table 1: Summary of representative tasks (§4.3) and their connections with different aspects of socially awareness (§2) including social factors, social interaction and social implication, as well as how these tasks should prioritize considerations that we have discussed in §4.1 for building socially aware NLP.

measures on static benchmarks or qualitative metrics like interviews. While such benchmarks show that many current systems lack simpler aspects of social awareness (Choi et al., 2023), relying solely on metrics omits the complexity and nuance of social awareness. Evaluation in the wild can help develop socially aware NLP that facilitates more human-AI interaction evaluation paradigms, and lead to new evaluation principles and protocols. (5) *Building socially aware language technologies for real-world uses* Socially aware language technologies are in a strong position to serve the diverse needs of users and contribute positively to society. One can leverage socially aware language technologies for positive impact (Jin et al., 2021) to create inclusive technologies by providing access to information and services for people with disabilities (Guo et al., 2020) or those who speak minority languages (Ziems et al., 2022a), as well as building technologies for crisis and emergency responses (Alharbi and Lee, 2022; Imran et al., 2016). (6) *Understanding how socially aware language technologies affect people and society* This essential area includes but is not limited to how such technologies affect how people communicate and interact with each other (Liu et al., 2022b); how such systems reinforce stereotypes or biases (Dev et al., 2022) and affect public trust, education, and the labor market (Eloundou et al., 2023); and how these technologies inform policy and regulation.

While not exhaustive, the above list sheds light on a few key directions for socially aware language technologies research. Some of these task categories, such as (2) and (6), have seen an increase in research lately, while others, such as (4) and (5), are still in early stages. Because more and more work is being done or needs to be done on social and language technologies, there is a crucial need for a sub-field of “*socially aware language technologies*”. Within this new sub-field, we need to ensure that

language processing advances are both technically sophisticated and socially aware. A unified sub-field focused on this goal would enable researchers to systematically address the challenges of embedding social intelligence into language models and facilitate more precise communication among scientists, policymakers, and the public. Therefore, the formal recognition of “*socially aware language technologies*” is a strategic step towards a future in which language technology responsibly intersects with human society.

5 Historical View of Socially Aware NLP

In the early days of AI, social awareness was assumed in the definitions, if not always explicitly stated (Turing, 1950; McCarthy et al., 2006). Early AI’s goal was to produce *human-like* behavior, which would inevitably include a degree of social awareness and, as a result, a tighter coupling of different aspects and disciplines. AI was initially conceived in a much more holistic manner than the fragmented space suggests today. Moravec’s paradox (Moravec, 1988), often summarized pithily as, “*In AI, easy things are hard, and hard things are easy*” (Pinker, 2003) has singled out *social awareness* and *motion* as the main areas where AI models have difficulties matching human performance even on simple tasks (while outperforming humans on tasks that require patience or logic). Over time, NLP and other subfields of AI focused on more easily solvable tasks, which usually meant purely information-based or logical tasks that did not require any social awareness. As a consequence, NLP has spent long stretches focusing on information-heavy linguistic tasks. Only recently have social, cultural, and demographic aspects of language seen a resurgence in research (Hovy and Yang, 2021; Dev et al., 2023). However, social awareness is difficult to define and implement.

The strong performance of LLMs on a variety of language understanding tasks could initially sug-

gest that these models also have social awareness. However, many of these tasks focus on language-only problems that do not require social awareness (i.e., that could be solved without recurrence to non-textual knowledge). Further, those tasks that purported to show social, psychological, or emotional aspects of models often operate under a flawed premise. For example, Sap et al. (2022) showed that while we can administer ToM tests to LLMs, that premise is flawed. A human subject’s ToM can be determined using a variety of question-based psychological tests because they answer as a result of their complex inner workings. LLMs, by contrast, respond by generating a list of likely words. Similarly, Shu et al. (2024) show that while LLMs can generate answers to psychometric questionnaires like personality tests, their answers are inconsistent and have little awareness of the premise. Thus, while the responses of humans and models are similar, they arise from very different causes. The development of these capabilities in the absence of explicitly programming will likely require new advancements that integrate social awareness in a responsible way—and corresponding new forms of evaluation that rigorously measure social awareness.

As we have reached the era in NLP dominated by LLMs, the next logical step is to tackle “harder” problems. Applying Moravec’s paradox, the next harder area for NLP would involve either motion (less applicable) or emotional intelligence and social awareness. This step coincides with a growing societal need. Yet, making progress in this area requires us to begin answering hard questions: Is social awareness gained gradually and/or systematically? Can we teach our machines how humans learn social awareness? In spite of the difficulty of replicating human social awareness in machines, we call for developing NLP systems that can learn and recognize social awareness over time, and to respond to these cues in a more human-like manner.

6 The Future of (Socially Aware) NLP

As LLMs take a more central role in AI research more broadly, many traditional NLP tasks have become obsolete. However, as the information processing power of those models grows, we are increasingly free to think about their use in a technological environment (Blodgett et al., 2020; Tedeschi et al., 2023; Abercrombie et al., 2023). We are not human because we speak; we speak because we are human. We are more than just language

factories, and language plays just one part in our complex social interactions. Language models, on the other hand, are language factories: capable of producing and processing words at astonishing rates, but lacking the faculties that drive human language production and processing.

Linguistics has long seen its subject (language) isolated from all other cognitive (and physical) capabilities. While that enables the study of certain aspects in isolation and develop theories and models, it obscured the larger picture. Sociolinguistics, psycholinguistics, and other sub-fields have fought hard to reintroduce these aspects to the larger linguistic field. NLP today is following a similar trajectory. With language models covering the basis for language production and comprehension, we can return to social aspects of language models. Understanding the social aspects of language technologies requires a refocus on emotional intelligence, cultural factors, values, and norms, as well as social interaction and the broader social implications. In addition, the future of socially aware NLP should emphasize ethical considerations and the responsible development process.

Socially aware NLP will likely transform industries and societal function, as well as shaping the broader field of AI, including but not limited to audio, vision, and robotics, where social awareness can play an even more critical role. For instance, the integration of social awareness in robotics can facilitate the development of robots that can safely and effectively interact with humans (e.g., elder-care robots, service robots) and lead to the advancements in computer vision that enable systems to better interpret emotions (Mittal et al., 2020; Kwon et al., 2023), social interactions, and cultural contexts from visual data (Kruk et al., 2023; Achlioptas et al., 2021). Again, we must proceed with a keen awareness of ethics and risks (Barrett et al., 2023).

In the future, we can also explore how these models work as social agents, which social cues they read and understand, and which tasks requiring social awareness they can master. This will require new tasks, metrics, and approaches, which are strikingly different from the goals we have followed as a field thus far. Most of all, it will require a re-alignment of the currently fractured AI landscape: we will need to work across fields to integrate models of emotions, values, and cultures into the models we have. There are plenty of unexplored research areas waiting to be explored.

7 Limitations

Our work focuses on the call for socially aware language technologies and provides an overview of what socially aware NLP is and how to build it. In spite of providing a general framework, our work does not provide any empirical evidence in the form of performance comparisons or other forms of quantitative measures. Additionally, socially aware NLP is linked to a variety of research fields, such as sociolinguistics, semantics, pragmatics, human-computer interaction, and social sciences; our current work only discuss the differences between socially aware NLP and a few very similar directions like computational social science and human-centered NLP. Our position paper also focuses only on a few key concepts related to socially aware NLP; nevertheless, there are many keywords and related subfields, from trust to equity in socio-technical systems, to transparency and societal values. Furthermore, we have discussed different dimensions of social awareness especially around social factors, social interaction and social implications, but it still remains challenging to computationally categorize, evaluate and even visualize social awareness in a system. Finally, we provide a general framework for defining socially aware NLP, but not specific technical methodologies to enable social awareness across different domains, as this would require diverse design choices, which we leave to future work.

8 Ethical Considerations

Developing socially aware NLP systems might require further information from users, and sometimes users may not be fully aware of and consent to such interactions, resulting in increased privacy risks. There is a risk that socially aware, more responsive language technologies can be misused to manipulate and personalize content, spread misinformation, or even persuade people towards certain decisions and behaviors. It is very likely that socially aware NLP systems will become out of date over time or in other contexts and should be continuously monitored in order to ensure that they meet users' varied needs in order to function safely and fairly. NLP systems that are socially aware also pose a risk of users' over-reliance, which can degrade human abilities and skills. Thus, we argue that, building socially aware NLP requires serious ethical considerations and a responsible development process, as addressing biases, ensuring

privacy, and taking into account potential misuse becomes increasingly important in building more transparent and accountable language technologies.

References

- Gavin Abercrombie, Amanda Curry, Tanvi Dinkar, Verena Rieser, and Zeerak Talat. 2023. [Mirages](#). on [anthropomorphism in dialogue systems](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4776–4790, Singapore. Association for Computational Linguistics.
- Panos Achlioptas, Maks Ovsjanikov, Kilichbek Haydarov, Mohamed Elhoseiny, and Leonidas J Guibas. 2021. Artemis: Affective language for visual art. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11569–11579.
- Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pages 337–371. PMLR.
- Alaa Alharbi and Mark Lee. 2022. Classifying arabic crisis tweets using data selection and pre-trained language models. In *Proceedings of the 5th Workshop on Open-Source Arabic Corpora and Processing Tools with Shared Tasks on Qur'an QA and Fine-Grained Hate Speech Detection*, pages 71–78.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Clark Barrett, Brad Boyd, Elie Bursztein, Nicholas Carlini, Brad Chen, Jihye Choi, Amrita Roy Chowdhury, Mihai Christodorescu, Anupam Datta, Soheil Feizi, et al. 2023. Identifying and mitigating the security risks of generative ai. *Foundations and Trends® in Privacy and Security*, 6(1):1–52.
- Federico Bianchi, Mirac Suzgun, Giuseppe Attanasio, Paul Röttger, Dan Jurafsky, Tatsunori Hashimoto, and James Zou. 2023. Safety-tuned llamas: Lessons from improving the safety of large language models that follow instructions. *arXiv preprint arXiv:2309.07875*.
- Steven Bird and Dean Yibarbuk. 2024. [Centering the speech community](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 826–839, St. Julian's, Malta. Association for Computational Linguistics.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is](#)

1044	for education. <i>Learning and individual differences</i> ,	<i>for Computational Linguistics: NAACL 2022</i> , pages	1098
1045	103:102274.	241–252.	1099
1046	Bhushan Kotnis, Kiril Gashteovski, Julia Gasteringer,	Yihe Liu, Anushk Mittal, Diyi Yang, and Amy Bruck-	1100
1047	Giuseppe Serra, Francesco Alesiani, Timo Sztyler,	man. 2022b. Will ai console me when i lose my pet?	1101
1048	Ammar Shaker, Na Gong, Carolin Lawrence, and	understanding perceptions of ai-mediated email writ-	1102
1049	Zhao Xu. 2022. Human-centric research for nlp: To-	ing. In <i>Proceedings of the 2022 CHI conference on</i>	1103
1050	wards a definition and guiding questions.	<i>human factors in computing systems</i> , pages 1–13.	1104
1051	Austin C Kozlowski, Matt Taddy, and James A Evans.	Weicheng Ma, Henry Scheible, Brian Wang, Goutham	1105
1052	2018. The geometry of culture: Analyzing mean-	Veeramachaneni, Pratim Chowdhary, Alan Sun, An-	1106
1053	ing through word embeddings. <i>arXiv preprint</i>	drew Koulogeorge, Lili Wang, Diyi Yang, and	1107
1054	<i>arXiv:1803.09288.</i>	Soroush Vosoughi. 2023. Deciphering stereotypes in	1108
1055	Julia Kruk, Caleb Ziems, and Diyi Yang. 2023. Impres-	pre-trained language models. In <i>Proceedings of the</i>	1109
1056	sions: Visual semiotics and aesthetic impact under-	<i>2023 Conference on Empirical Methods in Natural</i>	1110
1057	standing. In <i>Proceedings of the 2023 Conference on</i>	<i>Language Processing</i> , pages 11328–11345.	1111
1058	<i>Empirical Methods in Natural Language Processing</i> ,	John McCarthy, Marvin L Minsky, Nathaniel Rochester,	1112
1059	pages 12273–12291.	and Claude E Shannon. 2006. A proposal for the	1113
1060	Minae Kwon, Hengyuan Hu, Vivek Myers, Siddharth	dartmouth summer research project on artificial intel-	1114
1061	Karamcheti, Anca Dragan, and Dorsa Sadigh. 2023.	ligence, august 31, 1955. <i>AI magazine</i> , 27(4):12–12.	1115
1062	Toward grounded social reasoning. <i>arXiv preprint</i>	Trisha Mittal, Pooja Guhan, Uttaran Bhattacharya, Ro-	1116
1063	<i>arXiv:2306.08651.</i>	han Chandra, Aniket Bera, and Dinesh Manocha.	1117
1064	Maria Knight Lapinski and Rajiv N Rimal. 2005. An	2020. Emoticon: Context-aware multimodal emotion	1118
1065	explication of social norms. <i>Communication theory</i> ,	recognition using frege’s principle. In <i>Proceedings</i>	1119
1066	15(2):127–147.	<i>of the IEEE/CVF Conference on Computer Vision</i>	1120
1067	Anne Lauscher, Federico Bianchi, Samuel Bowman,	<i>and Pattern Recognition</i> , pages 14234–14243.	1121
1068	and Dirk Hovy. 2022. Socioprobe: What, when, and	Philipp Mondorf and Barbara Plank. 2024. Comparing	1122
1069	where language models learn about sociodemograph-	inferential strategies of humans and large language	1123
1070	ics. <i>arXiv preprint arXiv:2211.04281.</i>	models in deductive reasoning.	1124
1071	David Lazer, Alex Pentland, Lada Adamic, Sinan Aral,	Hans Moravec. 1988. <i>Mind children: The future of</i>	1125
1072	Albert-László Barabási, Devon Brewer, Nicholas	<i>robot and human intelligence.</i> Harvard University	1126
1073	Christakis, Noshir Contractor, James Fowler, Myron	Press.	1127
1074	Gutmann, et al. 2009. Computational social science.	Dong Nguyen, Laura Rosseel, and Jack Grieve. 2021.	1128
1075	<i>Science</i> , 323(5915):721–723.	On learning and representing social meaning in nlp:	1129
1076	Matthew Le, Y-Lan Boureau, and Maximilian Nickel.	a sociolinguistic perspective. In <i>Proceedings of the</i>	1130
1077	2019. Revisiting the evaluation of theory of mind	<i>2021 Conference of the North American Chapter of</i>	1131
1078	through question answering. In <i>Proceedings of the</i>	<i>the Association for Computational Linguistics: Hu-</i>	1132
1079	<i>2019 Conference on Empirical Methods in Natu-</i>	<i>man Language Technologies</i> , pages 603–612.	1133
1080	<i>ral Language Processing and the 9th International</i>	Ashwin Paranjape and Christopher D Manning. 2021.	1134
1081	<i>Joint Conference on Natural Language Processing</i>	Human-like informative conversations: Better ac-	1135
1082	<i>(EMNLP-IJCNLP)</i> , pages 5872–5877.	knowledgements using conditional mutual informa-	1136
1083	Belinda Z. Li, Alex Tamkin, Noah Goodman, and Jacob	tion. In <i>Proceedings of the 2021 Conference of the</i>	1137
1084	Andreas. 2023. Eliciting human preferences with	<i>North American Chapter of the Association for Com-</i>	1138
1085	language models.	<i>putational Linguistics: Human Language Technolo-</i>	1139
1086	Robert Litschko, Max Müller-Eberstein, Rob Van	<i>gies</i> , pages 768–781.	1140
1087	Der Goot, Leon Weber, and Barbara Plank. 2023.	Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Mered-	1141
1088	Establishing trustworthiness: Rethinking tasks and	ith Ringel Morris, Percy Liang, and Michael S Bern-	1142
1089	model evaluation. <i>arXiv preprint arXiv:2310.05442.</i>	stein. 2023. Generative agents: Interactive simulacra	1143
1090	Ruibo Liu, Ruixin Yang, Chenyan Jia, Ge Zhang, Diyi	of human behavior. In <i>Proceedings of the 36th An-</i>	1144
1091	Yang, and Soroush Vosoughi. 2023. Training socially	<i>annual ACM Symposium on User Interface Software</i>	1145
1092	aligned language models on simulated social inter-	<i>and Technology</i> , pages 1–22.	1146
1093	actions. In <i>The Twelfth International Conference on</i>	Alex Pentland. 2005. Socially aware computation and	1147
1094	<i>Learning Representations.</i>	communication. In <i>Proceedings of the 7th interna-</i>	1148
1095	Ruibo Liu, Ge Zhang, Xinyu Feng, and Soroush	<i>tional conference on Multimodal interfaces</i> , pages	1149
1096	Vosoughi. 2022a. Aligning generative language mod-	199–199.	1150
1097	els with human values. In <i>Findings of the Association</i>		

1151	Ildikó Pilán, Laurent Prévot, Hendrik Buschmeier, and Pierre Lison. 2023. Conversational feedback in scripted versus spontaneous dialogues: A comparative analysis. <i>arXiv preprint arXiv:2309.15656</i> .	Ashish Sharma, Kevin Rushton, Inna Wanyin Lin, David Wadden, Khendra G Lucas, Adam S Miner, Theresa Nguyen, and Tim Althoff. 2023. Cognitive reframing of negative thoughts through human-language model interaction. <i>arXiv preprint arXiv:2305.02466</i> .	1206
1152			1207
1153			1208
1154			1209
1155	Steven Pinker. 2003. <i>The language instinct: How the mind creates language</i> . Penguin uK.		1210
1156			1211
1157	Kim Plunkett. 1997. Theories of early language acquisition. <i>Trends in cognitive sciences</i> , 1(4):146–153.	Bangzhao Shu, Lechen Zhang, Minje Choi, Lavinia Dunagan, Dallas Card, and David Jurgens. 2024. You don’t need a personality test to know these models are unreliable: Assessing the reliability of large language models on psychometric instruments. In <i>Proceedings of the 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics</i> .	1212
1158			1213
1159	Vinodkumar Prabhakaran, Ajita John, and Dorée D. Seligmann. 2013. Who had the upper hand? ranking participants of interactions based on their relative power . In <i>Proceedings of the Sixth International Joint Conference on Natural Language Processing</i> , pages 365–373, Nagoya, Japan. Asian Federation of Natural Language Processing.		1214
1160			1215
1161			1216
1162			1217
1163			1218
1164			1219
1165		Mark Snyder and William Ickes. 1985. Personality and social behavior. <i>Handbook of social psychology</i> , 2(3):883–947.	1220
1166	David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? <i>Behavioral and brain sciences</i> , 1(4):515–526.		1221
1167			1222
1168		Yueqi Song, Catherine Cui, Simran Khanuja, Pengfei Liu, Fahim Faisal, Alissa Ostapenko, Genta Indra Winata, Alham Fikri Aji, Samuel Cahyawijaya, Yulia Tsvetkov, et al. 2023. Globalbench: A benchmark for global progress in natural language processing. <i>arXiv preprint arXiv:2305.14716</i> .	1223
1169	Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Explain yourself! leveraging language models for commonsense reasoning. <i>arXiv preprint arXiv:1906.02361</i> .		1224
1170			1225
1171			1226
1172			1227
1173	Surangika Ranathunga and Nisansa de Silva. 2022. Some languages are more equal than others: Probing deeper into the linguistic disparity in the nlp world. <i>arXiv preprint arXiv:2210.08523</i> .	Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Mireshghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, et al. 2024. A roadmap to pluralistic alignment. <i>arXiv preprint arXiv:2402.05070</i> .	1228
1174			1229
1175			1230
1176			1231
1177	Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. The risk of racial bias in hate speech detection . In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 1668–1678, Florence, Italy. Association for Computational Linguistics.		1232
1178			1233
1179			1234
1180		Christian Stab and Iryna Gurevych. 2014. Identifying argumentative discourse structures in persuasive essays. In <i>Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 46–56.	1235
1181			1236
1182			1237
1183	Maarten Sap, Ronan Le Bras, Daniel Fried, and Yejin Choi. 2022. Neural theory-of-mind? on the limits of social intelligence in large LMs . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 3762–3780, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.		1238
1184			1239
1185		Simone Tedeschi, Johan Bos, Thierry Declerck, Jan Hajič, Daniel Hershcovich, Eduard Hovy, Alexander Koller, Simon Krek, Steven Schockaert, Rico Sennrich, Ekaterina Shutova, and Roberto Navigli. 2023. What’s the meaning of superhuman performance in today’s NLU? In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 12471–12491, Toronto, Canada. Association for Computational Linguistics.	1240
1186			1241
1187			1242
1188			1243
1189			1244
1190	Deven Santosh Shah, H. Andrew Schwartz, and Dirk Hovy. 2020. Predictive biases in natural language processing models: A conceptual framework and overview . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 5248–5264, Online. Association for Computational Linguistics.		1245
1191			1246
1192			1247
1193			1248
1194			1249
1195		Michael Tomasello. 2014. <i>A natural history of human thinking</i> . Harvard University Press.	1250
1196			1251
1197	Omar Shaikh, Valentino Chai, Michele J Gelfand, Diyi Yang, and Michael S Bernstein. 2023a. Rehearsal: Simulating conflict to teach conflict resolution. <i>arXiv preprint arXiv:2309.12309</i> .	Alan M. Turing. 1950. I.—COMPUTING MACHINERY AND INTELLIGENCE. <i>Mind</i> , LIX(236):433–460.	1252
1198			1253
1199			1254
1200		Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. Language models don’t always say what they think: Unfaithful explanations in chain-of-thought prompting .	1255
1201	Omar Shaikh, Kristina Gligorić, Ashna Khetan, Matthias Gerstgrasser, Diyi Yang, and Dan Jurafsky. 2023b. Grounding or guesswork? large language models are presumptive grounders. <i>arXiv preprint arXiv:2311.09144</i> .		1256
1202			1257
1203			1258
1204		Eva Vanmassenhove, Christian Hardmeier, and Andy Way. 2018. Getting gender right in neural machine translation . In <i>Proceedings of the 2018 Conference</i>	1259
1205			1260
			1261

1262			1317
1263			1318
1264			1319
1265			1320
1266			
1267			
1268			
1269			
1270			
1271			
1272			
1273			
1274			
1275			
1276			
1277			
1278			
1279			
1280			
1281			
1282			
1283			
1284			
1285			
1286			
1287			
1288			
1289			
1290			
1291			
1292			
1293			
1294			
1295			
1296			
1297			
1298			
1299			
1300			
1301			
1302			
1303			
1304			
1305			
1306			
1307			
1308			
1309			
1310			
1311			
1312			
1313			
1314			
1315			
1316			