Transcending Fixed Meanings: Exploring the Impact of Linguistic Relativism on Adaptive Language Models in Generative AI

Subhangi Namburi
University of Nottingham, UK

subhangidasnamburi@gmail.com

Abstract: This research paper aims to start a discourse exploring the impact of linguistic relativism on adaptive language models within the field of generative AI, challenging the traditional fixed-meaning approach to natural language processing (NLP). It argues for a shift towards more personalised AI systems that can adapt to individual users' language nuances, rather than relying solely on large datasets with predetermined meanings. The current NLP models, based on conventional semantics, assume that language has a stable, objective reality where words have universally accepted meanings. This approach limits AI's ability to understand and generate language that reflects personal or contextual variations. The paper debates that generative AI should move towards a model that embraces the fluidity and subjectivity of language, where meanings are not fixed but can change depending on the speaker's intent or the situational context. This would involve incorporating user-specific data and situational awareness into AI systems, enabling them to interpret not just the literal meanings of words but also the speaker's intentions and the circumstantial cues that may alter these meanings. Such an approach would lead to the development of AI systems that are more adaptive and sensitive to the nuances of personal expression and contextual interpretation. However, the paper also acknowledges the potential ethical challenges associated with this approach. If AI systems are designed to allow for fluid and personalized meanings, they could be manipulated to shape public discourse in ways that reflect the biases or intentions of their developers. This raises concerns about the potential misuse of AI in influencing perceptions and realities, particularly when the fluidity of language is taken to an extreme where communication becomes chaotic and ineffective. Ultimately, suggesting personalised language models offer significant potential for enhancing AI's ability to understand and generate human-like language, there is a need for a balance between individual linguistic creativity and the communal aspects of language that ensure effective communication. The paper concludes that integrating linguistic relativism into AI models could advance the theoretical understanding of language in AI, but it must be approached with caution to avoid undermining the stability and clarity essential for meaningful human interaction.

Keywords: Linguistic relativism, Generative AI, Adaptive language models, Ethical AI, Language learning applications

1. The Postmodern Reconstructionist Approach to Language Learning

This paper explores how does the integration of generative AI in language learning applications reflect and reshape cultural narratives in language acquisition. By starting a discourse, I seek to delineate both the potential and limitations of AI-driven tools in creating meaningful, contextually grounded language learning experiences.

The postmodern reconstructionist approach to language learning debates the goal of language learning being that of transitioning from a Eurocentric 'dominator' model of culture; towards an aesthetic of interconnectedness, social responsibility and cultural attunement. This reconstructionist view of language learning provides optimism for language education by recognising language as existing within the culture, rather than alongside culture. Language then becomes a catalyst of change. The evolutionary perspective on language learning debates how closely intertwined language and culture are as adaptive mechanisms for human survival and socialisation. Language is a catalyst in understanding the cultural, social and political issues individuals face. It aids in coordinating their actions and social bonds. On the other hand culture as a shared system of beliefs, values, practices and traditions provides a sense of identity and belongingness within a group. Cultural practices and beliefs are communicated through language, and language reinforces cultural norms and values transmitting knowledge across generations, allowing cultures to evolve and adapt to changing circumstances over time and survive in a constantly changing environment. Language learning (LL) helps individuals interact with an international community, and interaction as an approach leads to more effective language development. The personal and interactional functions of language help shape the projected identity of an individual, their opinions and preferences and development of social bonds.

The 21st century welcomed a stream of language learning applications (LLA) aiding the reconstructionist approach to LL globally. The acceptance and use of the LLA surged because of the ability to teach and learn languages inter-continentally. This intertwined relationship between language and culture serves as the foundation for developing innovative approaches to language learning, using generative AI being one (Namburi & Hopkins, 2023).

2. Limitations of Generative AI in Understanding Context

In the field of language learning, there has been an ontological shock because of the recent advent of LLMs and generative AI. Generative AI, in the form of chatbots has penetrated into language learning, including those used in Duolingo and Memrise; they are fundamentally limited by their lack of embodiment and contextual awareness. They rely solely on textual or spoken input, missing out on the essential non-verbal cues that play a significant role in human conversations.

Image 1,2 and 3. Use of Generative AI in Duolingo

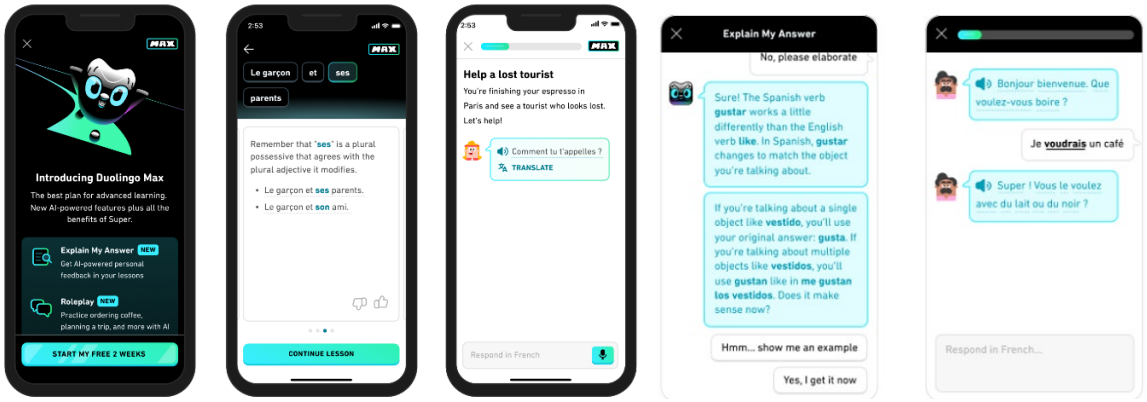

English Translation: Bot: Yes. Do you have pain in your eyes?

Human: No. I cry only on the inside.

Bot: Okay! So, what's up?

Image 4

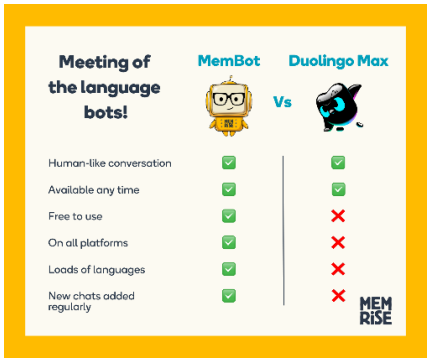

Memrise's claim of distinction over Duolingo in integrating AI into its software appears to be more a matter of marketing than substantive advancement. The benefits highlighted in their approach primarily reveal a superficial incorporation of AI, rather than a genuinely robust or transformative integration.

As a result, the chatbots' interactions are shallow and lack the depth needed for immersive language practice. They do not possess personal temporality or memory, which is crucial for understanding nuanced, long-term language learning. AI systems cannot retain long-term conversational context, limiting them to brief, turn-based interactions.

Use of language is not merely a process of symbol manipulation. Generative AI systems exist outside of the world, and only have access to data about the world that we provide to it. For this reason, the output of AI systems, even when they are linguistically competent, are ultimately mostly meaning-less and not the product of a true mind as we conceive it. Language and communication are highly dependent on context, and not just on general contexts like register, audience, etc. that could be given in a prompt to an LLM as an initial constraint. There are other contexts like the history you have with someone, previous conversations etc. LLMs are currently limited to a finite number of tokens that keep track of the context of a conversation, and this practical limit constitutes a memory context of about a few minutes in the real-world. We may gradually stretch these contextual limits as AI technology matures, but even under the best of scenarios it will be hard for AI agents to matures beyond a functionally limited conversational utility in the near future (Beals, 2024).

The problem with incorporating Generative AI in language learning applications is that AI relies on inductive reasoning and updates to generate responses. These systems are built on massive datasets and patterns but do not create new knowledge or form genuine explanations. Instead, they are restricted to rearranging pre-learned patterns. When you study and use a language for long enough, and make it a part of your own identity, it starts to change you. You adapt your thought processes to incorporate that language into your actual thinking. It becomes reflexive and natural. From a socio-linguistic standpoint the languages you acquire begin to inform your thinking and being as you navigate the world not just on a conscious level but also on the level of pre/unconscious thought. We develop shorthand modes of operation based on many assumptions of our world, including how we describe and explain it linguistically. Learning new languages is one way we have to develop new modes of being and understanding the world around us. Much of what human thought consists of on a moment-to-moment basis relies on pre/unconscious processes that, when functionally replicated by AI become highly resource-dependent and ungeneralisable to other domains.

This research paper hence aims to open a discourse about the practical implications and challenges of applying a flexible, personalised approach to language in the field of natural language processing (NLP) and generative AI and potentially influence both theoretical and practical advancements in the field. AI systems are designed to interpret and generate language based on fixed meanings derived from large datasets, which may not account for the nuances of personal or situational language use. This paper suggests a shift towards more personalised AI systems that can adapt to individual users' language nuances, challenging the current one-size-fits-all approach in AI.

3. AI Systems' Limitations in Language Acquisition

When AI chatbots generate sentences or content in LLAs, they do so based on statistical likelihood rather than a real understanding of context or user-specific learning needs. This limitation hinders the AI chatbots' ability to offer meaningful, adaptive learning experiences that go beyond surface-level interactions. Although chatbots simulate real-world conversations, the depth of understanding required to foster genuine language proficiency is still missing, making these interactions feel robotic and disconnected from the learner's actual progress or personal experiences. AI can now handle in-distribution data, but it still fails when encountering out-of-distribution situations, for example: when a learner brings up an unfamiliar cultural reference or uses language in a non-standard way, AI chatbots struggle to interpret or respond effectively. This makes the conversational practice less adaptive than it might initially appear, particularly for intermediate to advanced language learners who want to use language creatively.

AI systems cannot replicate this, limiting their utility in higher-level language acquisition. An additional concern also being the reliance on AI for learning can create biases and reinforce incorrect assumptions. This could manifest in reinforcing certain patterns or stereotypes embedded in the training data, leading to a less diverse or culturally sensitive learning experience. Furthermore, the lack of explainability in AI decisions—why a chatbot suggests one answer over another—can diminish trust and hinder the user's confidence in the learning process. These limitations of generative AI underscore a deeper issue regarding the understanding and flexibility of language, particularly as learners advance and introduce more creative, context-specific usage (Estes, 2022).

4. The Grounding Problem

Current AI systems, particularly Natural Language Understanding (NLU) models, grapple with the grounding problem, where symbols and language in AI do not inherently have meaning but rely on external interpretations. This critically raises the issue where symbols or words in a language model do not inherently have meaning. Instead, they rely on external interpretations—typically provided by humans or predefined rules—rather than having any intrinsic understanding of what the symbols represent. This becomes problematic when these models are expected to engage in tasks that require more than just pattern recognition or the mimicking of human language structures; they need to "understand" language in a meaningful way (Beals, 2024).

The symbols or words do not point to anything in the real world on their own. Instead, they are manipulated according to statistical patterns learned from large datasets of text. This is akin to handling words purely based on their form (e.g., the sequence of characters in a word) without understanding their underlying content or connection to the real-world phenomena they describe. This leads to the grounding problem—the model lacks a way to connect its internal symbols to the real-world objects, events, or experiences that those symbols are supposed to represent. Without grounding, AI cannot adapt its interpretation of language based on changes in personal context or nuance. Instead, it is locked into predefined associations, which restricts its flexibility and makes it unable to genuinely "understand" language as humans do. AI systems are then stuck at a surface level, manipulating symbols without grasping their deeper significance, thus making them inadequate for true conversational engagement (Estes, 2022).

AI relies on the use of relational knowledge bases, which, from a philosophical perspective, more closely resembles the associationist approach of conceptual structure (set theory). Whilst statistical approaches that train models on the intended result of a given query/command may be easier to train, and more generalisable to different applications, they face their own limitations. That is, they not only face the general challenges of statistical (deep) learning systems, but they are limited in their ability to interpret novel or ambiguous input—given their dependence on familiar examples from a given training set; as well as their lack of integration with real-world (background) knowledge; and commonsense reasoning (based on notions of causation, rather than mere correlation, between input and output).

From a philosophical perspective, AI's reliance on pre-existing patterns mirrors older conceptual frameworks, further limiting its ability to deal with ambiguity and novel contexts in real-time interactions. AI systems in LLAs should embrace linguistic relativism, wherein meaning is not fixed but instead varies according to context, personal intent, and social dynamics. Traditional NLP systems, with their reliance on fixed, objective meanings, inherently fail to capture this fluidity. The AI systems should be developed to interpret language based not just on static word meanings, but on the nuanced and evolving meanings that depend on personal and situational contexts. To overcome these limitations, AI systems should move toward embracing linguistic relativism, which allows for a more adaptable and nuanced understanding of meaning based on context.

To achieve this, AI systems would need to incorporate a deeper level of grounding—they would need to "understand" language not just in terms of statistical associations but in connection to real-world objects, events, and user-specific cues. This would allow AI to generate more context-aware responses and adapt to the speaker's unique way of using language. AI could further personalise its responses, offering more relevant and meaningful interpretations based on prior interactions; requiring AI to go beyond statistical associations and incorporate real-world knowledge and cues, making its responses more adaptable and contextually aware (Namburi & Hopkins, 2023).

Human testimony is a significant source of authoritative knowledge. AI systems processing and generating personalised language should incorporate individual testimonies; recognising testimonial authority, especially because AI knowledge sources are unavailable, highlighting the importance of context-driven testimony in shaping reliable language use. This notion of linguistic expression is also the source of authoritative knowledge when other more basic sources of knowledge like sense perception and inference are not available. The emphasis on testimony and trustworthiness in developing NLP models is crucial because it adds a layer to AI's potential for understanding language, showing how context-driven knowledge remains crucial (Language and testimony in classical Indian philosophy, 2020).

AI needs to employ a cognitive-functional grammar (also called a usage-based grammar). Language learners do not simply acquire abstract syntactic rules but rather grasp rules in both syntactic and semantic terms, considering patterns of use and communicative function. The current one-size-fits-all AI model overlooks these

crucial nuances, making it necessary to reconsider how AI could evolve to better capture personal and cultural diversity in language.

5. How Personalisation in AI Language Learning Misses the Mark on Cultural Sensitivity

The GPT model, and other transformer-based architectures have started incorporating context and generate more nuanced interpretations of language by considering the surrounding textual context. They do not operate purely on fixed, objective meanings anymore; instead, they adapt meaning dynamically based on the surrounding words and phrases. While these systems (LLA AI Chatbots) handle context better than older NLP models, they still fall short in terms of personal and cultural contextuality. Their understanding remains surface-level, unable to integrate deeply with personal histories, shared experiences, and non-verbal communication cues that are crucial for inclusive, accurate interactions. One of the critical limitations lies in their inability to maintain a continuous, deep understanding of long-term conversational history. These models typically operate within a finite token window, “remembering" a certain number of words or sentences from the current interaction. Once this limit is exceeded, earlier parts of the conversation are essentially forgotten, preventing the system from drawing upon previous interactions in any meaningful way.

AI is still unable to build on past conversations to form a deeper, more cohesive relationship with the user. This limitation hinders language learning, where the progress and learning trajectory of the user depend heavily on continuous, evolving interactions that accumulate knowledge over time. In human conversations, we adjust our speech and responses based on a shared history of interactions—a quality that is currently beyond the reach of most AI models.While memory-augmented models and retrieval-based methods are emerging, where previous conversations are indexed and recalled when relevant, this technology is still in its infancy. Even these systems struggle to distinguish between what information is contextually important to retain and what can be discarded, leading to either overly repetitive interactions or a lack of meaningful recall.

Language is not just a set of symbols and rules, but a cultural practice that reflects shared meanings, social norms, and values. When humans communicate, they draw not only on linguistic knowledge but also on a deep understanding of cultural references, idiomatic expressions, social expectations, and even historical or political context. A sentence can carry vastly different meanings depending on the speaker’s background, their relationship to the listener, and the specific cultural context in which the conversation takes place.

This gap is especially significant when considering language learning applications, where learners are not just trying to acquire vocabulary or grammar, but also to understand how a particular language is used in real-world, culturally new settings. In language learning, a failure to grasp these cultural nuances can lead to misunderstandings, miscommunication, or a superficial learning experience that doesn’t equip the learner to navigate complex social interactions in a foreign language. AI systems, when lacking this cultural sensitivity, risk reinforcing dominant cultural narratives while marginalising or misrepresenting minority cultures. This creates an incomplete and often biased representation of the language being taught.

Our language is influenced by our motivations, past experiences, current emotional states, and specific needs. The absence of personal temporality means that each interaction is treated in isolation, with little regard for the learner's progress over time or their individual learning style.

Language learning apps are trying to address this through personalisation techniques, such as user-specific profiles that store information about the learner's progress. However, these systems still fall short of the adaptive, intuitive teaching methods used by human educators who can infer meaning from subtle cues, respond to emotional states, and build on personal histories to create a more tailored and empathetic learning experience. Without these abilities, AI interactions remain transactional and impersonal, lacking the emotional depth that fosters long-term engagement and motivation in learners.

The reliance on statistical correlation rather than true understanding presents a profound philosophical challenge. These models are not "thinking" in the way humans do; they are merely calculating the most statistically probable sequence of words given a particular input. This lack of true understanding becomes evident when the AI encounters novel or ambiguous situations, particularly when the language used deviates from standard patterns or when cultural or emotional factors play a key role in shaping the meaning. One of my primary concerns regarding the rapid transformation of AI is that if there is meaning, then there are ideas; if there are ideas, then there is thinking; and if there is thinking, there is at least some degree of consciousness, or at least the beginnings of consciousness forming.

In human language learning, much of what we learn is shaped by abstract reasoning, inference, and the ability to create new meanings from limited information. We don't just learn language through repetition but also by developing conceptual understandings that allow us to generalise across different contexts. AI systems, in contrast, remain limited to the scope of their training data. They still struggle to interpret or generate language in ways that go beyond their statistical foundations, making it difficult for them to navigate ambiguity, handle metaphor, or understand language in its more creative and expressive forms.

On the technological front, AI systems would need to evolve to handle personal data efficiently by employing more compact models that can run on individual devices, reducing the reliance on cloud-based solutions and making real-time personalisation feasible. Advances in edge computing and privacy-preserving algorithms (like differential privacy) are crucial for balancing these demands. Progress in token-based memory management could enable AI to handle longer conversation histories without excessive computational overhead, improving the fluidity and continuity of interactions.

6. Navigating the Ethical Risks

The ethical risks of personalised AI systems, in language learning, are significant, systems could exploit personalised AI to influence or manipulate user opinions by tailoring the information or responses provided based on user-specific vulnerabilities or preferences. These systems could also reinforce linguistic or cultural biases by preferentially presenting dominant cultural narratives, while marginalising minority languages or cultural perspectives. Reinforcing the ethical framework in AI design would require safeguards like bias audits, transparent AI development processes, and multicultural training data that accounts for diverse linguistic and cultural perspectives. AI systems could employ a two-tiered language processing mechanism. The first tier would focus on ensuring communicative clarity by adhering to shared language norms that are widely understood within a community or language group. The second tier could allow for personalisation, where the system adapts to individual linguistic styles, preferences, and creativity.

AI could prioritise clarity by relying on communal linguistic standards as a baseline, while personal deviations would be contextually assessed. The system could flag when a personalised phrase is likely to cause confusion or misunderstanding, and either prompt the user for clarification or offer suggestions to rephrase while maintaining their stylistic intent. AI systems could learn over time how to balance these two aspects by tracking conversational success, ensuring that users are able to creatively express themselves without creating frequent communication problems.

7. AI's Role as a Supportive Tool, not a Creator of new Knowledge

The philosophical aspects of generative AI and it's use deserves greater attention and current discourse often neglects AIs identity of "knowledge creation". This deficit is because of philosophical misconceptions about the growth of knowledge generated using generative AI and the mistaken theoretical understand of generative AI knowledge production. This deficit is combined with philosophical misconceptions about the growth of knowledge using generative AI. This ignorance can give rise to problems in science and society that's why it is imperative to point out the flaws in our existing knowledge.

The potential of generative AI is what makes it even more important to have sufficient awareness of how current Generative AI works, what it can and cannot do. Generative AI works on the ideas of probability to a statement which is in line with an observation. Knowledge grows exactly the other way round. Observation does not precede theory, theory precedes observation (Velthoven & Marcus, 2024). Knowledge grows both in an individual mind as well as in society. AI only engages in calculation without engaging in deeper questioning or explanations. AI cannot advance science as it cannot explain why certain theories should be preferred over others. AI can nevertheless play an important supportive role in science, but AI cannot formulate new scientific explanations (Ordinary Language Philosophy, 2020). Current AI is built on principles on inductivism therein it cannot create new 'universal' or 'probable' statements based on individual instances. Generative AI systems give people the wrong impression that what AI does is similar to or even identical to how humans generate new knowledge.

In the field of generative AI, every individual instance results in an adjustment of the probability distribution. This means that specific instances impact the predictions for future events. The AI model's performance generally depends on the data we provide it with. In contrast, people acquire knowledge by virtue of explanations, not by 'extrapolating' from a collection of thousands of anecdotes. Therefore, AI should be considered as an instrument rather than a creator of new explanatory knowledge. Generative AI models

optimise their training task whenever they create samples that are close to the distribution of the training dataset. This is still entirely different than the way in which humans generate new ideas. When people generate new ideas, their new ideas may be completely unlike any idea that existed before.

The incorporation of generative AI into applications should be built with guardrails in place to address misconceptions and maintain intersectional accountability, recognising that the operation of AI is not akin to human thinking. Treating AI as ''just'' a technology also implies that the ultimate responsibility to explain the applications of AI should remain with people, not with the AI itself. A failure to do so could result in various adverse consequences, such as a wrong use of technology, lack of accountability and transparency, discrimination, and a mismanagement of expectations. Specifically, a misunderstanding of the underlying mechanism of knowledge creation can result in poor regulation and governance strategies (Velthoven & Marcus, 2024).

8. Reflections

In the rapidly evolving landscape of education, language learning is deeply intertwined with culture and identity. For centuries, the teaching and learning of languages has not merely been about acquiring vocabulary or mastering grammar but about immersing oneself in the social, historical, and cultural dimensions of the language in question. Yet, with the recent surge of generative AI, technology that threatens to reduce language learning to a sterile, contextually limited, and ultimately hollow experience. Generative AI, while powerful, should not be included in language learning due to its philosophical, cognitive, and pedagogical limitations, which are rooted in a misunderstanding of the nature of language, culture, and knowledge creation. At its core, language learning is not just about stringing together words or processing linguistic input. It is a dynamic process that serves as a key to cultural understanding and human connection. Language exists within culture, not alongside it, and is a force that can catalyse social change.

Generative AI, on the other hand, is fundamentally detached from these human experiences. It operates in a vacuum, relying on textual inputs to generate outputs, devoid of the rich, nuanced context that real-life language learning requires. Language learning is not a matter of extrapolating from a dataset of phrases or sentences. It involves constructing mental models of the worlds in which languages operate, shaped by human experiences, emotions, and cultural interactions. AI, however, cannot create new explanatory knowledge; it merely optimises its outputs based on the data it has been trained on. AI cannot think in a language, it only mimics thought and predicting knowledge is not knowledge. This philosophical deficit becomes especially apparent when considering the role of culture in language learning. Language and culture are adaptive mechanisms for human survival and socialisation. Cultural practices, beliefs, and values are transmitted through language, and these elements are not static, they evolve and adapt to new contexts over time. A language learner cannot simply acquire words and syntax; without immersing themselves in the cultural norms and values that those words encode. By reducing language learning to a transactional, probabilistic process, AI strips away the cultural richness that makes language learning a profound and transformative experience (Jenks, 2024).

The introduction of AI into language learning raises ethical concerns regarding transparency, accountability, and equity. AI systems are only as good as the data they are trained on, and if that data is biased or incomplete, the AI's outputs will reflect those deficiencies. This can lead to distorted representations of language and culture, reinforcing stereotypes or excluding certain voices from the language-learning process. Furthermore, the opaque nature of AI decision-making makes it difficult to hold these systems accountable when they fail to meet the needs of learners, potentially leading to a mismanagement of expectations and a lack of trust in the educational process. Moreover, the use of AI in language learning creates a false equivalence between human cognition and machine learning. AI systems give the impression that they can generate new knowledge in the same way humans do, but this is a dangerous misconception.

AI should be viewed as a tool rather than a replacement for human-led language education. It can at the most play a supportive role in certain aspects of language learning, such as vocabulary reinforcement or pronunciation practice, but it cannot replicate the deep, meaningful, and culturally immersive experience of learning a language in its true context. The ultimate responsibility for language education must remain with human educators, who can provide the cultural and contextual knowledge that AI systems lack. By relying on AI for language education, we risk reducing language to a sterile, transactional process that overlooks the deeper cultural and philosophical dimensions of learning.

The analogy drawn from a conversation between Humpty Dumpty and Alice in 'Through the Looking-Glass' highlighting the tension between fixed and fluid meanings of language, which mirrors my discourse analysis on the role of AI in language learning.

Image 5.

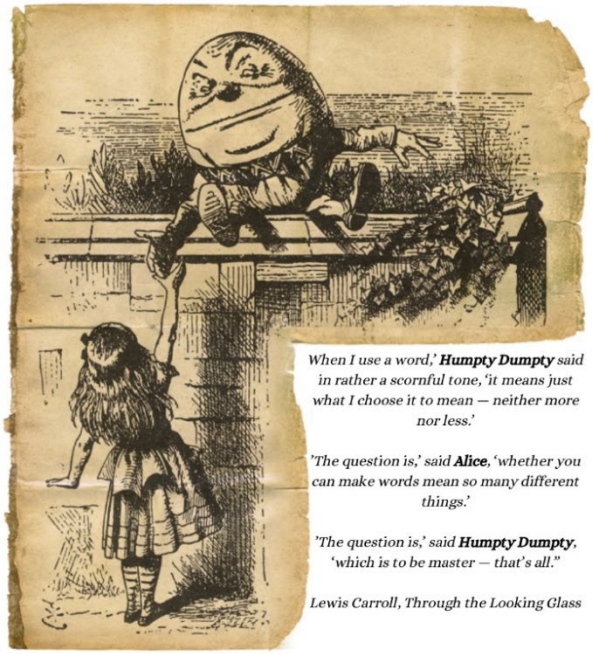

In this conversation, **Humpty Dumpty** asserts, "When I use a word, it means just what I choose it to mean—neither more nor less." Highlighting his belief in the arbitrary nature of language, suggesting that words do not have inherent meanings but rather meanings that are assigned by us.

Alice, on the other hand, represents a conventional view of language, wherein words have specific meanings that are universally accepted and understood; assuming an objective reality to the meaning of word.

This is in direct contrast to **Humpty Dumpty**'s subjective and individualistic approach. **Humpty Dumpty**'s perspective is essentially about the relationship between words and their meanings. Instead of words having a fixed meaning and us learning about what they signify; suggesting that the individual should be allowed to dictate what words mean, bending language to personal will and context.

This philosophy challenges the conventional semantics and destabilises the common belief that language is a shared resource with agreed-upon meanings.

Image 6.

Humpty Dumpty presents a nuanced understanding of naming and language exploring the fluidity and malleability of linguistic significance from a distinctly personal perspective.

Humpty Dumpty's perspective on language, represents a form of linguistic relativism, where the meanings of words are not fixed but are subject to individual interpretation.

Alice then contrasts this perspective with a traditional, semantic view that sees language as a set of shared symbols with common meanings that is used for standardised communication among people.

Humpty Dumpty advocates for a radical form of authorial control over language where the speaker has the ultimate authority over the use of words suggesting a personalised somewhat solipsistic relationship with language where intent is in foreground independent of conventional or societal norms.

Alice's is a normative stance where words have stable, consistent meanings that are agreed upon within a linguistic community and is used as a tool for achieving understanding and cooperation between individuals.

Humpty Dumpty's approach challenges the idea of language being a transparent medium for expressing truth or reality. Instead, it suggests that language is a manipulable construct, shaped and reshaped by users.

In postmodernist and constructivist theories about language the role of language is seen as a tool for constructing reality rather than merely describing it. NLP models of AI are built on a **Alice** like perspective where language, words have a fixed meaning that are understood and processed based on statistical regularities in data.

Image 7.

Language is a tool for constructing realities rather than merely describing them. Generative AI does not just interpret input(words) but actively constructs outputs that influence perceptions and realities.

The ability of AI to generate content that may redefine terms or concepts in new and unforeseen ways parallels **Humpty Dumpty**'s notion of linguistic control.

Generative AI models need to be built on **Humpty Dumpty**'s model of language wherein language can be contextualised, personalised and meaning can change based on the speakers intent or the situational context.

Generative AI models need to be built on **Humpty Dumpty**'s model of language wherein language can be contextualised, personalised and meaning can change based on the speakers intent or the situational context.

To be able to do so AI would have to not only determine the literal meaning of the words but also understand the speakers intentions, circumstantial cues that might even alter the literal meaning. This could be done by incorporating user specific data and situational awareness. Embracing the **Humpty Dumpty**-like approach to language would lead to the development of more personalised AI systems that adapt to individual users' language and meanings. These systems would then need to become adaptive and sensitive to the nuances of personal expression and contextual interpretation.

A potential ethical problem could be that because meaning is fluid and controlled by individuals (programmers of AI systems) AI could be misused in shaping public discourse. How the AI models are trained, the data inputted, and the intention of the developer could influence how these systems understand and use language, potentially leading to biases, misunderstandings, or manipulations of information.

Image 8.

Humpty Dumpty's perspective empowers the individual to use language creatively, allowing for personal expression and the capacity to influence and redefine terms within one's own communicative context. It acknowledges the fluid and dynamic nature of language, reflecting how meanings can evolve over time or in different contexts, aligning with the postmodernist understandings of linguistics and semantic change.

Nonetheless there is potential threat in hindering effective communication. If every human being could assign personal meanings to words, mutual understanding could be jeopardised, making communication chaotic and ineffective.

Since it ignores social conventions and shared meanings, **Humpty Dumpty**'s approach could be seen as an extreme form of linguistic relativism. It neglects the communal and social aspects of language that facilitate coordinated and cohesive social interaction. Its practicality is also in question when considering shared concepts such as law, science, and education. There is an arbitrary nature of the linguistic signs which AI does not understand. It does not understand the connection between the signifier (sound pattern) and the signified (concept) that is arbitrary and socially constructed. Mutability and immutability of language signs, which evolves through social use but retains some level of arbitrariness. These signs are not fixed but shaped by context and speaker intent. The idea of language evolving yet constrained by tradition inevitably creates balance between linguistic creativity and the need for stable, shared communication. The system of language is inherently relational—where meanings are defined not by intrinsic value but by their differences from other signs—needing AI to move beyond conventional semantics. AI models, under Saussure's perspective, would then need to understand language as a network of shifting relations rather than a collection of static definitions (Saussure, 1959).

Alice represents a neo colonial, consensus-based view of language, where words have fixed, stable meanings agreed upon by a linguistic community that threatens the existence of non- American cultures and usage of these languages. Her perspective aligns with conventional semantics, seeing language as a shared system that facilitates understanding and cooperation among people.

Image 9.

Alice's perspective ensures clarity and efficiency in communication, relying on shared definitions that are globally understood, reducing misunderstandings and increasing effective communication; hence playing a crucial role in the identity and cohesion of communities. But this could at the same time restrain linguistic evolution, potentially stifling the natural dynamism of language as it evolves with societal changes. It would also exclude or marginalise non-standard dialects and the innovative use of language, which can be important for artistic and personal expression. By strictly adhering to agreed-upon meanings, it limits how individuals can express unique perspectives or adapt language to new contexts and ideas.

A middle ground between **Humpty Dumpty** and **Alice** could be one where we acknowledge that while individuals can influence language, the community ultimately plays a critical role in stabilising meanings through consensus and usage patterns.

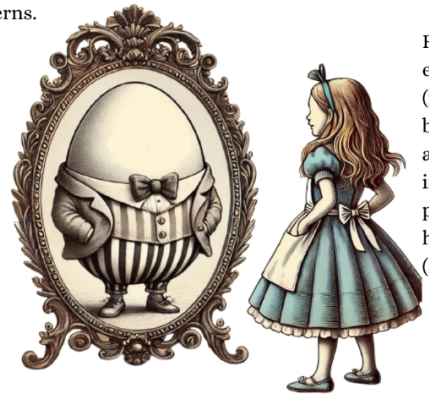

Furthermore, modern linguistic theories, proposed by Saussure, emphasises the arbitrary nature of signifiers (words) and their signifieds (meanings). From this perspective, **Humpty Dumpty**'s approach may not be as radical as it initially appears; it simply takes the inherent arbitrariness of language to its logical extreme by personalising it. By incorporating flexibility into AI models, this research pushes for a more personalised approach to language processing, aligning it closer to how humans generate meaning and build relationships through language (Saussure, 1959).

Acknowledgements

This paper is lovingly dedicated to my mother, who sat with me every day as a child, patiently writing down the words I invented. She wrote down each word, along with the meaning I had assigned to them, and then rewrote new ones every day, as I forgot everyday what I had created the day before. She made me believe that my thoughts, my language, and my words held immense power, even at the age of five. Her unwavering faith in my ability to create and the fact that she wrote them all down in a notebook is what has stayed with me till date. Thank you for celebrating each new word and meaning, and for never questioning their place in our conversations. I also want to thank my father for never questioning why his daughter was using hundred new words every day, treating it as the most natural thing in the world. Normalising it is what shaped my unrelenting trust in my ability to create.

References

Beals, J. (2024). Grounding AI: Understanding the implications of generative AI in world language & culture education.

Estes, G. A. (2022). Linguistic flexibility: How a more fluid approach to classroom language can help students succeed.

Jenks, C. J. (2024). Communicating the cultural other: trust and bias in generative AI and large language models.

Language and testimony in classical Indian philosophy. (2020). Stanford Encyclopaedia of Philosophy. [Online]. Available at: https://plato.stanford.edu/entries/language-india/

Namburi, S. & Hopkins, G. (2023). Meraki: Encouraging Language Learning in Real-World Simulations through AI Para-Social Relationship Building

Ordinary Language Philosophy. (2020). Internet Encyclopaedia of Philosophy. [Online]. Available at: https://iep.utm.edu/ord-lang/

O'Regan, J. P. & Ferri, G. (2024). Artificial intelligence and depth ontology: implications for intercultural ethics.

Saussure, F. de. (1959). Course in general linguistics. Edited by C. Bally and A. Sechehaye. Translated by W. Baskin.

Velthoven, M. J. & Marcus, E. J. (2024). Problems in AI, their roots in philosophy, and implications for science and society.