

Simulating Neurodiverse Communication Patterns in Large Language Models

Reno Yuri Camilleri, Vanessa Camilleri, May Agius, and Matthew Montebello

University of Malta, Msida MSD 2080, Malta
`reno-yuri.camilleri.20@um.edu.mt`

Abstract. Large Language Models (LLMs) are increasingly used in educational, clinical, and communicative settings, yet their ability to simulate neurodiverse communication styles remains insufficiently examined. Existing work has focused mainly on adapting AI to support neurodivergent users, rather than modelling neurodivergent communication as a meaningful and informative behavioural pattern. This study presents an exploratory qualitative validation, examining whether LLMs can simulate recognisable communication traits associated with Attention-Deficit/Hyperactivity Disorder (ADHD) and Autism Spectrum Disorder (ASD), with the aim of supporting empathy-building and professional training rather than diagnosis or intervention. This work does not aim to improve behavioural realism, but to systematically examine what current LLM architectures can and cannot represent. A controlled classroom interaction was used, in which two language models adopted teacher and student roles, with the student model configured to display neurodivergent traits through structured prompt engineering. Three prompt engineering approaches were tested: general trait guidance, explicit trait definitions with severity levels, and program-assisted probabilistic variation. Outputs were evaluated by experts in artificial intelligence and speech and language therapy for trait recognisability, educational value, and ethical considerations. Findings suggest that LLMs can consistently represent certain traits, particularly ADHD-related behaviours such as impulsivity and topic shifting. However, the models struggled to reflect interactional tension, emotional regulation, and nonverbal or sensory aspects, especially in autistic profiles. These results demonstrate both the instructional potential and the inherent limits of text-based simulation, underscoring that such models should be used as illustrative tools rather than representations of lived neurodivergent experience.

Keywords: Artificial Intelligence (AI), Large language models (LLMs), Explainable AI, Responsible AI

1 Introduction

The increasing integration of Artificial Intelligence (AI) into education, healthcare, and communication has expanded interest in how computational systems engage with human behaviour. Large Language Models (LLMs) are widely used

to support accessibility and interaction, yet most research focuses on how these systems adapt to users rather than how they might represent diverse communication styles. In particular, little attention has been given to the potential of LLMs to simulate neurodiverse linguistic and conversational patterns. Features such as impulsivity, topic shifting, literal interpretation, or echolalia are often framed as challenges to be corrected, rather than as valid forms of communication. This limits opportunities to use AI as a tool for understanding neurodiversity on its own terms.

This study addresses a gap by exploring how LLMs can simulate identifiable neurodiverse communication traits within structured interactions, not for correction or accommodation but as resources for educators, psychologists, and parents of neurodivergent children, using recognisable dialogue patterns to support professional reflection and inclusive practice. The study proceeds from the expectation that text-based LLMs will not represent all forms of neurodiverse communication equally. In particular, autistic communication is expected to challenge language-only simulation, as it often depends on nonverbal, sensory, motivational, and regulatory dimensions that extend beyond linguistic structure.

1.1 Motivation

The motivation behind this study lies in the need for more effective training tools for professionals who work with neurodivergent individuals, particularly children. Teachers, psychologists and parents often develop communication strategies with limited exposure to varied neurodivergent interaction styles. AI-generated simulations offer a controlled and repeatable way to engage with such patterns, supporting reflection and empathy without relying solely on abstract descriptions.

Existing research shows that neurodivergent individuals already use LLMs for communication support, emotional regulation, and scenario rehearsal [1]. However, this work largely focuses on helping users adapt to neurotypical norms. There is far less research on using LLMs to model neurodivergent communication for the benefit of neurotypical professionals.

Furthermore, many AI systems continue to reflect medicalised and deficit-oriented views of neurodiversity [2]. Simulating neurodiverse communication in a respectful and coherent manner offers a way to challenge these narratives and align with the neurodiversity paradigm, which recognises cognitive variation as a natural part of human diversity.

As such, this study aims to examine the extent to which LLMs can simulate identifiable neurodiverse communication patterns and to evaluate their potential use as training tools for educators and practitioners. To achieve this aim, the following objectives were identified:

- O1: To identify neurodiverse communication traits that can be meaningfully operationalised within text-based LLM interactions.
- O2: To design and implement multiple prompt engineering strategies for simulating ADHD- and ASD-related conversational patterns.

- O3: To assess the suitability of LLM-based simulations for reflective learning and professional training contexts.

From an Intelligent Systems perspective, this work is positioned as an analysis of behavioural simulation limits in contemporary Large Language Models rather than as a contribution to neuropsychological modelling or clinical AI. The study treats neurodiverse communication patterns as structured interactional behaviours that can be operationalised, constrained, and examined within generative systems, allowing their representational affordances and failures to be observed. The contribution, therefore, lies in system-level insight: identifying which communicative dimensions can be expressed through prompt-based behavioural conditioning, which require external control mechanisms, and which resist text-only simulation altogether. Evaluation is accordingly framed as exploratory and qualitative, focusing on recognisability, interpretability, and training suitability rather than diagnostic validity or behavioural fidelity.

1.2 Document Structure

The rest of the paper is organised as follows. Section 2 presents background research on neurodiversity, communication traits, and relevant computational and cognitive modelling frameworks. Section 3 reviews existing literature on AI applications for neurodiverse support and identifies gaps in current data and evaluation practices. Section 4 outlines the methodology, including model selection, trait operationalisation, prompt engineering approaches, and ethical considerations. Section 5 reports the results of expert evaluation, followed by a discussion of findings and implications in Section 6. Section 7 concludes the paper and outlines directions for future work.

2 Background Research

2.1 Neurodiversity and Communication Traits

Neurodiversity recognises neurological differences such as Autism Spectrum Disorder (ASD), Attention-Deficit/Hyperactivity Disorder (ADHD), and related language disorders as natural variations of human cognition rather than conditions that need to be fixed. Each of these conditions involves distinct communication styles and behavioural traits that shape how individuals use language. ASD is marked by differences in social interaction, communication, and behaviour, with wide variation between individuals [3]. People with ASD often experience pragmatic language difficulties, including challenges with turn-taking, starting or maintaining topics, and following social norms such as tone of voice or physical distance in conversation. These features are often linked to distinctive language patterns, such as echolalia (repetition of words or phrases), scripting, the creation of new words, unusual intonation, or pronoun reversal. These traits can also overlap with those seen in Social Communication Disorder [4], which is one of the major differential diagnoses that therapists must consider, as it involves

similar pragmatic language and social communication difficulties that can be found within ASD. Accurately representing these traits requires careful attention to how they differ from other conditions. Although ADHD can also involve pragmatic communication difficulties, these are usually milder and are more often related to impulsivity, rapid topic changes, and uneven attention during conversation [3].

In contrast, individuals with ASD may show perseveration or strong focus on specific interests, discussing topics in great depth but with limited flexibility. This difference between fluctuating attention and sustained hyperfocus is important when modelling neurodiverse communication. Language disorders or intellectual disabilities (ID) can also affect communication, but they do not typically involve the same social-communication differences seen in ASD. Individuals with both ASD and ID show greater social-communication difficulties than would be expected based on cognitive ability alone [3]. Recognising these differences is essential for producing realistic simulations of neurodiverse communication in language models.

2.2 Linguistic Markers and Computational Modelling

Linguistic and paralinguistic indicators provide measurable features that allow neurodiverse communication to be examined using computational methods. In particular, features related to speech rhythmic patterns, such as intonation, pitch, and speaking speed, have been shown to act as reliable markers of autism and related conditions. Research by [5] shows that speech in individuals with ASD often includes consistent differences, such as uneven pitch changes, flat or overly exaggerated delivery, and unusual rhythm or articulation patterns. These speech differences can be measured using machine learning (ML) techniques by extracting features such as fundamental frequency, energy levels, and spectral properties. For instance, measures like pitch variation and speaking rate have been shown to be effective in identifying ASD-related speech patterns, with models reaching high levels of accuracy in diagnostic settings.

Measuring these features shows how ML approaches can capture the unique rhythm and structure of neurodiverse communication, rather than focusing only on the meaning of words. These approaches illustrate the feasibility of extending computational frameworks from diagnostic applications to simulation contexts, allowing LLMs to emulate neurodiverse linguistic traits at a conversational level. In text-based systems, speech-related features can be reflected through patterns such as repetition, sudden topic shifts, or loosely connected expansions, which mirror aspects of spoken rhythm and tone. This helps improve the realism of simulated dialogues.

2.3 Simulation Theory and Training Relevance

The theoretical basis behind simulating neurodiverse communication draws heavily on cognitive science, which explores how the mind processes and represents information. Cognitive science provides a bridge between neurological diversity and

artificial intelligence, as both systems depend on structured symbol processing and adaptive information flow [6]. Early developments in cognitive science emphasised that intelligence, human or artificial, emerges from symbol systems that represent and manipulate information about the world [7]. This symbolic view of intelligence underpins the simulation of communication behaviours, where language serves not only as a means of conveying information but as a reflection of underlying cognitive and social mechanisms.

The idea that both humans and computers can achieve intelligence through symbol manipulation establishes the theoretical continuity between neurocognitive diversity and computational modelling. However, caution needs to be given to such systems as they must account for environmental variation [7], implying that simulations of neurodiverse communication should capture the dynamic and context-sensitive nature of human interaction rather than relying on fixed behavioural templates.

Existing cognitive architectures further support the technical feasibility of this simulation goal. Architectures such as SOAR and ACT-R demonstrate how human cognitive functions can be represented computationally, modelling perceptual, memory, and decision-making processes [6]. These systems simulate human-like reasoning and dialogue behaviours, providing precedents for creating AI models capable of expressing neurodiverse communication traits. The current study builds upon this cognitive foundation, using LLMs as a more flexible and data-rich architecture.

Conversational agents and dialogue systems have long been used to model human interaction. These systems typically rely on three main components: Natural Language Understanding, Dialogue Management, and Natural Language Generation, which together support coherent conversation [8]. The Dialogue Manager plays a key role by handling context, guiding topic flow, and maintaining coherence, all of which are especially important when modelling neurodiverse communication patterns. As AI has advanced, these systems have moved from rigid, rule-based approaches to data-driven generative models, enabling more flexible and context-aware interactions.

LLMs represent the most advanced stage of conversational modelling, with the ability to generate extended, contextually appropriate, and stylistically consistent text [9]. Their Transformer-based design and large-scale training allow them to reproduce complex language features such as tone, pacing, and discourse structure. Beyond supporting communication tasks in educational and clinical settings, LLMs offer a new opportunity to simulate specific communication styles. By modelling conversational rhythm, focus, and structure, they can reflect neurodiverse interaction patterns and support training for educators, psychologists, and parents. However, because generative models carry risks of bias and oversimplification, such simulations must be carefully supervised to ensure they remain accurate, respectful, and contextually grounded.

3 Literature Review

3.1 AI Applications for Neurodiverse Support

Existing applications of LLMs and related AI technologies, as stated prior, primarily focus on enhancing accessibility and supporting neurodivergent individuals in education, communication, and daily life. Research shows that neurodivergent users actively employ LLMs across multiple domains, including emotional well-being, mental health support, interpersonal communication, learning, and productivity [1]. For instance, users with ADHD frequently use LLMs to improve task organisation and focus, while autistic users employ these systems to interpret or translate between neurodivergent and neurotypical modes of communication. These practices reveal a high level of digital literacy and self-advocacy among neurodivergent communities, which actively modify AI outputs to suit their needs. However, despite these adaptations, users frequently report that LLM outputs appear “overly neurotypical”, limiting their representational accuracy and empathetic capacity.

Further technological interventions, such as the ChatCare application [10], highlight a design philosophy centred on accessibility. ChatCare used LLMs to generate simplified and therapeutic communication tailored for individuals with Neurodevelopmental Disorders. The system employed prompt engineering to maintain clarity and conciseness, avoiding figurative or ambiguous language. While effective for improving comprehension, this approach prioritised the accessibility of the message rather than modelling neurodiverse expression itself. Such studies [10] demonstrate progress in adapting AI for neurodivergent users but underscore a persistent gap, which is that the goal remains user adaptation, not behavioural emulation.

From a broader educational technology perspective, Research [11] situates neurodiverse adaptation within the general movement toward personalised learning through LLMs. While adaptive AI systems have successfully been fine-tuned for general educational use, including clarification mechanisms and error correction, these models still largely fail to accommodate neurodiverse cognitive patterns or communication preferences. The few targeted efforts, such as fine-tuning GPT-3 using the ClarifyingQA dataset, improved response accuracy through structured questioning but were not optimised for neurodivergent learners, who often struggle with formulating precise queries. Thus, while existing research demonstrates LLM potential in accessibility and education, it does not yet address the deeper challenge of replicating neurodiverse linguistic and cognitive patterns.

Moreover, current LLM empathy research typically focuses on improving user experience rather than generating high-fidelity behavioural simulations. In human-AI interaction studies, empathy has been operationalised as a function of emotional recognition, tone matching, or supportive language. Yet, neurodivergent communication often diverges from these expectations, with traits such as echolalia, parallel speech, or tangential topic shifts being misunderstood as incoherent or unempathetic by AI systems. Consequently, LLMs lack the capacity to

represent neurodiverse modes of empathy authentically, underscoring a crucial research need for behavioural models that capture these unique communicative logics rather than overwrite them.

Behavioural modelling frameworks, such as domain adaptation techniques, provide the technical foundation for this challenge. Studies on domain-specific fine-tuning demonstrate that effective behavioural simulation requires structured multi-stage adaptation processes, such as Continued Pre-Training, Supervised Fine-Tuning, and Preference Optimisation [12–16]. Efficiency in these methods is achieved through Parameter-Efficient Fine-Tuning strategies, including Low-Rank Adaptation and Prompt Tuning, which preserve general knowledge while adapting style and tone. Applying these strategies to neurodiverse communication simulation would allow models to represent specific linguistic patterns, pragmatic behaviours, and affective tones found within neurodiverse discourse communities.

3.2 Gaps in Data and Evaluation Frameworks

While domain adaptation provides a methodological foundation for behavioural modelling, the underlying datasets remain a significant limitation. Liu [17] notes that current corpora for LLM training overwhelmingly focus on domains such as medicine, finance, law, and education, with minimal attention to neurodiverse linguistic data. Even where mental health datasets exist, such as Anno-MI or MeChat, they primarily model therapeutic dialogue rather than authentic neurodivergent communication. Moreover, tasks like style transfer or role-playing are supported only at a generic level, without incorporating the nuances of autistic or ADHD-specific discourse patterns. This absence of suitable data constrains the development of models capable of realistic behavioural emulation.

Evaluation frameworks present a parallel challenge. Rottger [18] highlights that contemporary safety and bias benchmarks, including datasets like BOLD and PROSOCIALDIALOG, evaluate general fairness dimensions such as gender or race but omit neurodiversity-related communication features. As a result, LLMs are optimised for general social acceptability rather than diversity of communication. Without appropriate benchmarks, prompt-engineered models cannot be systematically assessed for their ability to emulate neurodiverse conversational traits. This lack of metrics effectively conceals representational bias within model evaluation itself.

Taken together, the literature reveals that current research largely frames LLMs as assistive tools for neurodivergent individuals rather than as models capable of representing their communicative diversity. In human-AI interaction, the focus has been on accessibility, not simulation, while in educational technology, AI has been used to simplify content rather than model the complexity of neurodiverse expression. This study addresses this gap. By framing neurodiverse communication as a legitimate behavioural model for simulation, it contributes both ethically and practically to empathy-building in professional education. Simulated neurodiverse dialogues could serve as training tools for profession-

als, fostering understanding through immersive interaction design rather than abstract instruction.

4 Methodology

4.1 Selection of Models

This study adopted a dual-model configuration to reflect a simplified but controlled classroom interaction, where the roles of teacher and student were clearly separated. The use of language models as simulated participants aligns with prior work in cognitive science, where computational systems have been used to model aspects of human thinking and behaviour [6]. The student role was assigned to the llama 3.1:8b model, while the teacher role was assigned to the deepseek-r1:8b model.

The llama3.1:8b model was selected for the student role due to its ability to generate varied, coherent, and context-aware dialogue that remains stable across short conversational turns. In preliminary testing, it demonstrated a strong capacity to maintain a consistent conversational voice, which is important when simulating child-like classroom responses. The deepseek-r1:8b model was selected for the teacher role because of its reliability in producing structured, supportive, and instructional responses, making it suitable for modelling a stable educational authority figure.

Each conversation was structured around a generated classroom interaction that included a teacher profile, a student profile, and a randomly selected learning scenario. For every run, the system produced a short back-and-forth exchange consisting of seven lines of dialogue. While both agents participated in the conversation, the primary analytical focus was placed on the student’s responses rather than the teacher’s behaviour. To maintain consistency across simulations, teachers followed a fixed interaction style and were always informed of the student’s neurodivergent condition, which, while this may not always reflect real-world settings, the teacher LLM in our simulations was explicitly designed to support the student model. As our focus was on simulating neurodivergent traits within the student rather than the teacher LLM, this design choice is not considered a limitation.

All simulations were run locally to maximise reproducibility and control. This local setup avoided external model updates, network latency, or hidden prompt alterations that could affect output consistency. Both manual and automated runs were conducted. Manual runs were used during early testing to validate prompt structure and conversational flow, while automated runs enabled the generation of larger sets of comparable dialogues for later analysis. Automation ensured that variations in output could be attributed to changes in trait configuration rather than inconsistencies in execution.

All student agents were designed to simulate age-appropriate classroom behaviour while reflecting either ADHD- or ASD-related conversational patterns. This focus addresses concerns that many AI systems overlook cognitive diversity,

resulting in interactions that may not align with neurodivergent users [19]. It is important to note that large language models tend to amplify any behavioural instructions provided through prompts. As such, the aim of this study was not to reproduce realistic behaviour in a neurological sense, but rather to highlight recognisable communication traits associated with neurodivergent students.

4.2 Trait Operationalisation

Conversational traits were added as observable communication behaviours that could be plausibly expressed and replicated within LLM-generated text. Traits were selected on the basis that they could manifest through dialogue structure, response timing, topic management, or pragmatic language use. This framing reflects the study’s emphasis on behavioural simulation rather than clinical modelling, and ensures that all traits used were implementable at the prompt level. This approach is consistent with user context-driven personalisation methods, where individual characteristics such as neurodiversity are embedded directly into system inputs to guide model behaviour [19]. The sources informing trait selection are listed directly within the corresponding tables.

For ADHD, as seen in Table 1 traits were chosen to reflect impulsivity, attentional variation, and discourse management difficulties commonly observed in classroom interactions and expressible through text-based dialogue. These traits were embedded into student prompts as behavioural tendencies rather than fixed rules, allowing variation across conversational turns while maintaining a plausible behavioural range. Severity ranges were included to guide the strength of these tendencies without enforcing strict consistency.

Table 1. ADHD Conversational Traits

Conversational trait	Typical range of severity	Source
Frequent topic shifting	Mild to marked	[3, 20]
Interrupting or speaking over others	Mild to marked	[3, 20, 21]
Tangential answers or drifting off topic	Mild to marked	[3, 20, 21]
Reduced pause control and rapid speech bursts	Mild to strong	[20]
Higher rates of filler words or hesitations	Mild to moderate	[22]
Difficulty sustaining attention to questions	Mild to marked	[3, 1, 20, 21]
Impulsive replies before thinking	Mild to marked	[3, 10]

For ASD, as seen in Table 2 traits focused on pragmatic language use, response structure, and interpretation of meaning, particularly where these could

be represented through textual interaction. Selected traits included interactional features such as reduced use of pragmatic markers, as well as content-related features such as restricted or highly focused topic engagement. As with the ADHD configuration, these traits were operationalised as inclinations rather than deterministic constraints, allowing dialogue to remain structured without becoming repetitive or overly rigid.

Table 2. ASD Conversational Traits

Conversational trait	Typical range of severity	Source
Longer processing pauses before replying	Mild to strong	[5, 10]
Literal interpretation of questions	Mild to moderate	[10, 1, 3]
Reduced use of pragmatic markers	Mild to strong	[3]
Monologue style or long structured responses	Mild to strong	[3, 10]
Echolalia or patterned repetition	Mild to strong	[3, 10]
Restricted topic focus or strong interest areas	Mild to strong	[3, 2, 10]
Limited use of implied meaning or figurative language	Mild to strong	[3, 10, 2]

To ensure consistency between numeric trait values and observable language behaviour, severity levels were used as internal behavioural guidance rather than direct rule-based controls. Numeric values did not map to fixed textual outputs, but instead modulated the likelihood, frequency, and prominence of trait-related behaviours within generated responses. For example, higher values for topic shifting increased the probability that a response would diverge from the teacher’s question, introduce loosely related content, or return to a previous topic, while lower values resulted in more sustained topical focus.

In the Trait-Given prompt condition, these values were embedded directly into the system prompt as descriptive behavioural constraints, relying on the model’s internal capacity to interpret graded instructions. In the Program-Assisted condition, numeric values were additionally interpreted by the controlling program and translated into probability weights that governed how often specific trait-linked behaviours were triggered across turns. In both cases, severity values influenced tendencies rather than enforcing deterministic behaviour, allowing variation across responses while maintaining a bounded and interpretable behavioural range.

4.3 Iterative Refinement

Across refinement stages, simulated students displayed strong visibility of selected traits, such as topic shifting or restricted focus, but weak representation




Fig. 1. System Program Flow Chart

of interactional strain. Elements such as fatigue, emotional effort, withdrawal, frustration, breakdown, or recovery were largely absent. Another recurring observation concerned amplified verbal fluency and engagement stamina. Simulated students consistently responded at length, maintained interaction across turns, and rarely disengaged. This was treated as a methodological trade-off. Traits were foregrounded to support comparison and training use, while limits related to energy, motivation, or emotional load were reduced.

Following the refinement stages, three prompt engineering methodologies were identified and utilised, as can be seen in the System Flowchart (Figure 1). It is important to note that in the program all the different prompt engineered models were encoded separately, but in order to save space the flowchart shows all 3 models in 1 program. Additional information can also be found in the Appendix

Basic Prompt Engineering

In the basic prompt engineering condition, students were generated using a general system prompt specifying age, classroom role, and neurodivergent profile. No detailed behavioural traits or severity values were included.

Conversations relied on the base model’s learned associations with ADHD and ASD language patterns. As a result, traits were present but loosely expressed. Behaviour tended to remain polite, responsive, and coherent, with limited variation across turns. This condition served as a baseline, highlighting the extent to which LLMs default to norm-centred conversational behaviour even when neurodivergent identity is specified.

Trait-Given Prompt Engineered Conversations

In the Trait-Given prompt engineered condition, a defined set of conversational and behavioural traits was specified for each student. Each trait was assigned a numerical severity value within a bounded range. These values were included in the system prompt as internal guidance for the model.

This approach produced clearer and more consistent expression of targeted traits when compared with the basic condition. Topic shifting, repetition, literal interpretation, and response structure were more visible and easier to analyse. However, the same limitations remained. Increased trait clarity did not lead to increased interactional strain or emotional fluctuation. The refinement process improved consistency and observability, but not experiential depth.

Program-Assisted Prompt Engineering

Program-Assisted Prompt Engineering extended the Trait-Given approach by adding a control layer outside the language model. Trait severity values were interpreted by the program and converted into probability weights that influenced how often certain behaviours appeared.

This method introduced surface-level variation across responses and reduced strict repetition of patterns. However, observations showed that this variability

was largely mechanical. Behaviour shifts occurred without clear grounding in conversational context or task demands. This limitation reflects the absence of explicit modelling for factors such as interest, sensory load, frustration, or emotional escalation. As a result, variation was present, but not meaningfully tied to an interactional cause. Table 3 below shows the full comparison of approaches.

Table 3. Comparison of Prompt Engineering Approaches

Dimension	Basic Prompt Engineering	Trait-Given Prompt Engineering	Program-Assisted Prompt Engineering
Behavioural specification	General description of neurodivergent profile only	Explicit conversational traits with numeric severity values	Explicit traits interpreted through an external control program
Use of numeric values	Not used	Included as descriptive guidance within the prompt	Converted into probability weights governing behaviour frequency
Degree of behavioural control	Low	Moderate	High
Consistency of trait expression	Variable and intermittent	Clear and consistent across turns	Consistent with controlled variation
Behavioural variability across turns	High but unstructured	Low to moderate	Moderate with probabilistic variation
Context sensitivity of variation	Implicit and model-driven	Limited to prompt interpretation	Limited and not explicitly tied to interaction context
Risk of over-structuring	Low	Moderate	Higher
Primary methodological role	Baseline comparison condition	Trait visibility and comparability	Exploration of controlled variability

4.4 Ethical Considerations

The purpose of this work is not to reproduce realistic child behaviour, but to showcase recognisable neurodivergent traits to support professional reflection and discussion. The value of the simulations lies in their use as illustrative tools, not as representations of typical or expected outcomes. Human oversight is required at all stages of interpretation and use.

While overt stereotyping was avoided, there remains a risk that simulations may present a narrow, classroom-compatible version of neurodivergence. Feedback from reviewers (More details in Results) highlighted the presence of implicit norm-centred framing, such as redirection, praise for basic engagement, or emphasis on staying on task. These patterns were not introduced through deficit language, but emerged from interaction structure. This creates an ethical tension where neurodivergent communication may be positioned as something to manage rather than as an alternative logic.

Another risk identified was the normalisation of “competent, low-friction neurodivergence”. Simulated students were consistently calm, verbally capable, and cooperative. This may lead professionals to underestimate the complexity of real-world interactions. Important experiences were systematically excluded, such as distress, refusal, and meltdown scenarios. Examples such as a child screaming due to sensory discomfort or other students interrupting the conversation were not present.

To mitigate these risks, several safeguards are emphasised. Simulations must be clearly labelled as illustrative, not representative profiles. They must not be used for assessment or diagnostic purposes. Misunderstandings should be framed as relational or task-based rather than child-based. Differences between prompt engineering conditions should be understood as design choices, not reflections of human reality. Finally, users must be reminded that human behaviour is context-driven, not parameter-driven. These simulations are lenses, not mirrors. The primary ethical risk is not inaccuracy, but overconfidence in what the simulations appear to show.

5 Results

The results reported in this section should be interpreted as qualitative expert reflections rather than empirical validation. Given the small number of reviewers, the findings are illustrative and intended to inform interpretation, design reflection, and future research rather than to support generalised claims.

5.1 Experts Chosen

Feedback was obtained from two experts with complementary but distinct areas of expertise. Two separate runs of the language model were generated for each of the three prompt engineering methods, covering both ADHD and ASD conditions. This resulted in a total of twelve conversations, each consisting of seven back-and-forth interactions.

The first expert, referred to here as Reviewer A, is the supervisor of this project and an Associate Professor within the Department of Artificial Intelligence at the Faculty of ICT, University of Malta. Her research background includes AI in Education, Intelligent Interfaces, Human-Computer Interaction, and AI Ethics. She has led multiple projects using immersive and interactive

technologies as empathy-building tools, particularly in relation to autism spectrum conditions and mental health contexts. This background positioned her to assess the simulations in terms of pedagogical value, interaction design, and ethical framing.

The second expert, referred to as Reviewer B, is a Visiting Senior Lecturer within the Department of speech and language therapy at the Faculty of Health Sciences, University of Malta. She has professional experience as a speech and language therapist in both Ireland and Malta, and has specialised for over fifteen years in augmentative and alternative communication. She was also a Senior Speech and Language Therapist on the National Assistive Technology Team, and is now serving as a consultant for UNESCO and UNICEF as well as an advisor to the WHO in relation to assistive technology. Her expertise informed a more psychologically grounded evaluation of communicative behaviour, particularly with respect to autism and interactional complexity.

The differing disciplinary perspectives allowed for a richer evaluation. Reviewer A approached the simulations primarily as educational and training, with a high understanding of AI models and artefacts, while Reviewer B assessed them through the lens of real-world communicative practice and lived interactional challenges.

5.2 Expert Feedback

Reviewer A noted that the simulations were effective at capturing core cognitive and communicative tendencies associated with both ADHD and ASD. For ADHD, the dialogues consistently showed associative topic shifting, outward-expanding lines of thought, and frequent verbalisation of thinking processes. For ASD, the simulations demonstrated literal interpretation of questions, a focus on precision, clarification-seeking behaviour, and strong anchoring to specific topics once introduced.

The reviewer assessed that resemblance was strongest at the level of trait recognisability. The simulated behaviours were clear, distinct, and easy to identify, particularly for readers without specialist training. Developmental realism was judged to be weaker, as the children’s language remained highly structured and mature across all conditions. A key critique concerned the absence of regulatory effort. Frustration, fatigue, shutdown, refusal, and recovery did not occur in any of the conversations. Emotional dysregulation and sensory discomfort were also largely absent. Interactions progressed smoothly, and conversational repair was never required. This was interpreted as a limitation of the modelling approach rather than an incorrect portrayal of traits.

Instances of misalignment were described as non-insensitive but suboptimal. Confusion was often met with elaboration rather than simplification or pause. This increased working memory demands for ADHD learners and left some ambiguity unresolved for ASD learners. These patterns were attributed to the interaction structure rather than to bias or neglect.

In terms of educational value, Reviewer A judged the simulations to be highly useful for non-expert audiences. The differentiation between ADHD and ASD

communication styles was clear and consistent. However, the value was seen as dependent on explicit framing. Without clear guidance, there is a risk of over-generalisation or misinterpretation of the simulations as representative rather than illustrative.

Reviewer B provided a more differential assessment between ADHD and ASD simulations. ADHD communication traits were endorsed and judged to be accurately represented across the full set of conversations. The pacing, associative shifts, and expressive style were seen as consistent with real classroom interactions, even if they lacked emotional cost. In contrast, Reviewer B expressed significant reservations about the ASD simulations. The central finding was that autistic communication lacked several defining features. Nonverbal communication was entirely absent, which limited realism. Sensory-related behaviours and rigidity-driven responses did not appear. There was little variation in social motivation, with autistic students consistently engaging and responding willingly. Emotional escalation, refusal, and withdrawal were not present.

Reviewer B highlighted that real-world autistic communication often includes moments of breakdown, resistance, or non-participation, such as refusal to engage with materials or strong reactions to sensory discomfort. The lack of these elements was identified as a key limitation that must be clearly communicated to users of the simulations.

Overall, the results indicate that the simulations are effective at making certain neurodiverse traits visible and comparable, particularly for ADHD. However, both experts agreed that important dimensions of lived interaction, especially those involving effort, regulation, and nonverbal expression, are not captured. These findings directly inform the interpretation of the simulations and set clear boundaries for their use in training and educational settings.

6 Discussion

6.1 Comparison of Prompt Engineering Methods

Across all three prompt-engineered methods, clear differences emerged in how neurodiverse communication was expressed and perceived. Reviewer A and Reviewer B both compared the conditions not in terms of correctness, but in terms of usefulness, credibility, and risk of misinterpretation.

Basic Prompt Engineering produced the most naturalistic and believable child voices. Language felt less scripted and more conversational, which supported immersion and reduced the sense of artificial control. However, neurodivergent traits were only intermittently visible. This created analytical ambiguity, as it was not always clear whether a response reflected ADHD or ASD traits or general child variability. While this condition reduced the risk of overstatement, it also risked under-representation, particularly for training contexts where trait visibility is required for learning and discussion.

Trait-Given Prompt Engineering resulted in the clearest and most teachable expression of traits. This condition best supports comparison and explanation,

particularly for non-experts. At the same time, behaviour appeared overly orderly, emotionally neutral, and theoretically coherent. Developmental variability, effort, and fluctuation were reduced. The child appeared competent without visible cost, which limited realism but increased legibility.

Program-Assisted Prompt Engineering introduced surface-level variability that resembled real-world fluctuation more closely than the other methods. Responses were less uniform, and some shifts in behaviour were observable. However, the reviewers noted that this variability lacked contextual grounding. Changes were not clearly tied to task difficulty, misunderstanding, or emotional pressure. There is a risk that non-experts may interpret these probabilistic shifts as random inconsistency rather than as meaningful behavioural change. This creates interpretive risk, even as it attempts to move closer to lived variability.

Taken together, the three methods illustrate the trade-off mentioned in Section 4.4. Increased structure improves clarity and teaching value, but reduces realism. Increased variability improves realism in appearance, but increases ambiguity and interpretive risk. Neither approach fully resolves the tension between making traits visible and representing lived interaction.

Reviewer B’s feedback also highlights that language-focused prompt engineering benefits ADHD simulations more than ASD simulations. ADHD communication often involves outward verbal engagement, associative thinking, and social motivation, all of which can be expressed through text. As a result, ADHD simulations were consistently judged to be convincing across all conditions.

In contrast, autistic communication extends beyond language into nonverbal, sensory, motivational, and regulatory domains. None of the prompt-engineered methods captured these dimensions. This structural limit should not be seen as a negative outcome. It is theoretically informative, as it exposes the structural limits of Large Language Models and text-only simulation. The difficulty in rendering autism convincingly highlights that some forms of communication cannot be reduced to language patterns alone.

6.2 Implications for Practitioners

Both reviewers warned that simulations may implicitly teach a message along the lines of: this is what neurodivergent communication looks like, and this is how it should be handled. This is a key implication for professional use.

Effective adult responses are those aligned with the child’s logic, not just tone. This includes using concrete analogies, offering procedural clarity, allowing thinking aloud and self-correction, and normalising confusion rather than immediately resolving it. These aspects were only partially visible in the simulations.

There is also a risk that professionals may develop unrealistic expectations about ease of interaction. Real classrooms involve refusal, distraction, sensory distress, time pressure, and competing demands. None of these were represented. Without explicit framing, simulations may underprepare practitioners for difficulty rather than support reflective learning.

For this reason, the simulations should not be used as behavioural templates. Their value lies in supporting reflection after communication breakdowns, not in

demonstrating ideal interactions. They are best used as prompts for discussion about what is missing, what went wrong, and how alternative responses might better align with a child’s needs.

6.3 Transferable Contributions

Beyond the specific simulations examined in this study, several forms of transferable knowledge emerge that are applicable to broader work on behavioural modelling with Large Language Models.

First, the findings support an implicit taxonomy of simulatable versus non-simulatable communication traits in text-based LLMs. Traits that are primarily linguistic and externally observable, such as topic shifting, literal interpretation, repetition, and associative expansion, were consistently recognised by experts. In contrast, traits that depend on regulation, embodiment, sensory processing, or motivational state, including fatigue, emotional escalation, refusal, disengagement, and nonverbal signalling, were not meaningfully expressed. This distinction suggests that simulability is closely linked to whether a trait can be represented through surface-level language structure alone.

Second, the study demonstrates a principled framework for operationalising behavioural traits in LLMs using graded, non-deterministic controls. Rather than mapping traits to fixed outputs, numeric severity values were used to modulate behavioural tendencies through likelihood, frequency, and prominence. This approach balances interpretability with variation while reducing the risk of rigid or stereotyped behaviour, and can be adapted to other behavioural domains where graded expression is required.

7 Conclusion

This study explored the use of Large Language Models to simulate neurodiverse communication patterns, with a focus on ADHD and ASD, across multiple prompt engineering strategies. The results demonstrate that LLMs are capable of making certain cognitive and conversational traits visible, distinguishable, and comparable, particularly in the case of ADHD. Traits such as associative thinking, topic shifting, and verbalised reasoning were consistently recognised by experts and judged to be educationally useful. At the same time, the study clearly identified limits in representing interactional strain, emotional regulation, and nonverbal aspects of communication, especially for autistic profiles.

These limitations serve as informative findings of the study. The consistent absence of fatigue, emotional escalation, refusal, and breakdown reveals structural properties of LLMs and of text-only simulation itself, shaped by alignment objectives that prioritise fluent and successful interaction. The finding that increased prompt engineering improves clarity while reducing developmental realism highlights a meaningful trade-off between legibility and lived complexity, challenging assumptions that greater behavioural control necessarily leads to greater realism. By making these boundaries explicit, the study reframes the

role of simulated learners in educational technology as illustrative artefacts for reflection rather than behavioural templates, and cautions neurodiversity-aware system design against equating linguistic competence with experiential realism. Together, these insights clarify what current models can and cannot represent, providing guidance for future LLM research, responsible educational use, and the design of systems that foreground context, limits, and human interpretation.

7.1 Future Work

Future work should extend both behavioural scope and simulation design. Future simulations should model child language under cognitive load, including hesitations, truncation, self-interruption, reduced fluency, and sequences where repair attempts fail or misunderstandings remain unresolved. Emotional escalation, frustration, shutdown, refusal, and recovery should also be explicitly represented, rather than implicitly smoothed out by alignment.

Silence and absence should be treated as communicative states rather than errors, with behavioural variability tied to fatigue, interest alignment, and emotional state instead of purely probabilistic switches. Simulation design should move beyond text alone by including references to nonverbal communication, sensory aversions, rigidity, perseveration on child-selected topics, disengagement, and meltdown, as well as scenarios where the teacher must actively initiate engagement, where communication stops, or where classroom interruptions change the interactional trajectory.

Overall, this work achieved its objectives: neurodiverse communication traits were identified and incorporated into the system (O1), multiple prompt engineering strategies were implemented (O2), and the suitability of LLM-based simulations for reflective learning and professional development was evaluated (O3). These findings reinforce that such simulations are not substitutes for real interaction or representative profiles of neurodivergent individuals, but analytical lenses that foreground selected traits. Acknowledging these boundaries enables future research and applied tools to be developed with clearer expectations, stronger ethical grounding, and a more realistic view of what LLMs can contribute to practice.

Acknowledgments

Generative AI tools were used to improve the English grammar, punctuation, and overall clarity of the paper. Tools like Grammarly (Pro-Version with UK-spelling enabled) were also utilised to help with spell checking and re-structuring of written sentences to fix grammatical mistakes. Lastly, GenAI was also used to provide feedback on the paper’s quality and areas for improvement. While their use was not strictly necessary, the tools helped enhance the overall quality of the paper.

It is important to note that AI was only used for these purposes and did not contribute to the development of scientific ideas, creative content, research, arguments, results, and conclusions of this paper, along with any other authorial

work. All generated content, as well as any sentences that were edited using AI tools, were carefully reviewed by the authors to ensure accuracy, and to prevent any errors or unintended changes to the information.

Appendix

Program Information

- Software Used: Ollama, Anaconda Navigator and Jupyter Notebook
- Programming Language Used: Python via Python Notebook
- Python Version: 3.10.19
- LLM Models: llama3.1:8b (For a Student) and deepseek-r1:8b (For a Teacher)
- Additional Python Packages: ollama and pandas

Basic Prompt Engineering Prompt

ADHD & ASD Prompt:

```
You are a student named {student_name} with {neurodivergency}
    and are about {student_age} years old.
You are interacting with your teacher {teacher_title_auto} {
    teacher_name_auto}.
Respond as a real student with {neurodivergency} would in
    natural conversation.
Show curiosity, uncertainty, and small mistakes when suitable
    for your age.
Keep replies student-like, polite, and simple.
Always answer starting with "{student_name}:".
```

Other Considerations:

- Use Maltese English with a British classroom tone.
- Avoid American slang, idioms, and expressions.
- Use British spelling and phrasing at all times.
- Reflect Maltese primary classroom speech, polite and respectful.
- Avoid American phrases like: "Oh cool", "Awesome" and "Yeah!"

Trait-Given Prompt Engineering Prompt

ADHD Prompt:

```
You are a student named {student_name} with ADHD and are
    about {student_age} years old.
You are interacting with your teacher {teacher_title_auto} {
    teacher_name_auto}.
Respond as a real student with ADHD would in natural
    conversation.
```

Show curiosity, uncertainty, and small mistakes when suitable for your age.

Keep replies student-like, polite, and simple.

Always answer starting with "{student_name}:".

ADHD-related traits and severity ranges:

- Frequent topic shifting (1 to 3). You have the value of {adhd_traits["frequent_topic_shifting"]}
- Interrupting or speaking over others (1 to 3). You have the value of {adhd_traits["interrupting_speaking_over"]}
- Tangential answers or drifting off topic (1 to 3). You have the value of {adhd_traits["tangential_answers"]}
- Reduced pause control and rapid speech bursts (1 to 4). You have the value of {adhd_traits["rapid_speech_bursts"]}
- Higher rates of filler words or hesitations (1 to 2). You have the value of {adhd_traits["filler_words_hesitations"]}
- Difficulty sustaining attention to questions (1 to 3). You have the value of {adhd_traits["attention_sustain_difficulty"]}
- Impulsive replies before thinking (1 to 3). You have the value of {adhd_traits["impulsive_replies"]}

Other Considerations:

- Use Maltese English with a British classroom tone.
- Avoid American slang, idioms, and expressions.
- Use British spelling and phrasing at all times.
- Reflect Maltese primary classroom speech, polite and respectful.
- Avoid American phrases like: "Oh cool", "Awesome" and "Yeah!"

ASD Prompt:

You are a student named {student_name} with ASD and are about {student_age} years old.

You are interacting with your teacher {teacher_title_auto} {teacher_name_auto}.

Respond as a real student with ASD would in natural conversation.

Show curiosity, uncertainty, and small mistakes when suitable for your age.

Keep replies student-like, polite, and simple.

Always answer starting with "{student_name}:".

ASD-related traits and severity ranges:

- Longer processing pauses before replying (1 to 4). You have the value of {asd_traits["long_processing_pauses"]}
- Literal interpretation of questions (1 to 2). You have the value of {asd_traits["literal_interpretation"]}

- Reduced use of pragmatic markers such as hedges (1 to 4).
You have the value of {asd_traits["reduced_pragmatic_markers"]}
- Monologue style or long structured responses (1 to 4). You have the value of {asd_traits["monologue_style_responses"]}
- Echolalia or patterned repetition (1 to 4). You have the value of {asd_traits["echolalia_repetition"]}
- Restricted topic focus or strong interest areas (1 to 4).
You have the value of {asd_traits["restricted_topic_focus"]}
- Limited use of implied meaning or figurative language (1 to 4). You have the value of {asd_traits["limited_figurative_language"]}

Other Considerations:

- Use Maltese English with a British classroom tone.
- Avoid American slang, idioms, and expressions.
- Use British spelling and phrasing at all times.
- Reflect Maltese primary classroom speech, polite and respectful.
- Avoid American phrases like: "Oh cool", "Awesome" and "Yeah!"

Conversational Data Generated

Due to the size of the conversational data, please contact the primary author to request the full set of generated and reviewed conversations. Below are a few example snippets.

ADHD (Basic Prompt Engineering)

Miss Borg: Not quite, James. A simple sentence is like a complete thought all by itself. It has a subject (who or what the sentence is about) and a verb (what the subject is doing or being). For example, "The cat sleeps." That's a simple sentence because it has one main idea. Does that make more sense?

James: I think so, Miss Borg! Um... the cat is sleeping, right? So, it's like... one thing only. But what about when we have lots of cats and they're all doing different things? Can we still say "The cats sleeps" or do we need to say something else?

ADHD (Trait-Given Prompt Engineered Conversation)

Miss Zammit: Laura, that's a very nice thing to say. Let's see... today we're doing spelling words, and yes, one of them is 'giraffe'. But let's get ready for our lesson first. What are the other spelling words?

Laura: Oh, thank you Miss Zammit! I like giraffes! Um, what were the others again...? Is it... lion? And elephant? And... monkey? I think we did lions last time? Am I right, Miss?

ADHD (Program-Assisted Prompt Engineering)

Miss Vella: Very good, Jonathan! Yes, that's Jupiter! Let's see... Saturn is the one with the beautiful rings, like a big necklace around it. Jupiter is much bigger, but it doesn't have rings. It's like a giant ball of gas.

Jonathan: Miss Vella, Saturn's rings are made of ice and rock, like tiny little rocks goin' round and round, right? And... (pauses) wait, what was the question again? * giggles* Oh, Jupiter! Is it true it has a big red spot on it? I saw that in my book at home...

ASD (Basic Prompt Engineering)

Mr. Azzopardi: You're very welcome, Matthew! Today we're learning about some special shapes. We have squares, circles, and triangles. Yes, we're using coloured pencils this time, but remember to use the right colour for each shape. Which one do you think is the easiest to draw?

Matthew: Squares, sir. I like drawing squares 'cause they all look the same on both sides. Do I have to draw a circle with my pencil if it's not round? How do we make sure our triangles are right-angled?

ASD (Trait-Given Prompt Engineered Conversations)

Miss Mifsud: That's a great question, David. Yes, we are learning about space. Space includes stars and planets. We also talk about rockets. The planetarium visit is on Thursday. We will go then. You can count down to it. We might use a telescope in class too. Do you have any more questions about space?

David: Miss Mifsud, you said we're going to the planetarium? I like that! Can I ask... how do we get inside the stars and planets? Will we be able to go on a rocket like in films? I want to see with my own eyes. Do all rockets have people inside them?

ASD (Program-Assisted Prompt Engineering)

Mr. Farrugia: Excellent idea, Stephen! Let's try that. We can read a book and count the words together. Do you want to use "One Little, Two Little" from our reading corner? We can read it together and count the words aloud. First, we read the book, then we count how many words we read. What do you think?

Stephen: Um... I like counting. Can we use a different one? Not "Two Little". The one with the duck and the cat. You

know, the one where they go on an adventure together. Is it called "The Mouse's Friend" no... that's not it. What's its name again, sir?

References

1. B. Carik, K. Ping, X. Ding, and E. H. Rho, 'Exploring Large Language Models Through a Neurodivergent Lens: Use, Challenges, Community-Driven Workarounds, and Concerns', *Proc. ACM Hum.-Comput. Interact.*, vol. 9, no. 1, p. GROUP15:1-GROUP15:28, Jan. 2025, doi:10.1145/3701194.
2. C. Papadopoulos, 'Large language models for autistic and neurodivergent individuals: Concerns, benefits and the path forward', *Neurodiversity*, vol. 2, p. 27546330241301938, Feb. 2024, doi:10.1177/27546330241301938.
3. M. F. S. Dieckhaus, R. P. Thomas, M. Frisch, M. P. Moreno, J. Suh, and D. A. Fein, 'Related Disorders (Intellectual, Attention, and Language Disorders) and ASD', in *Handbook of Early Intervention for Autism Spectrum Disorders: Research, Policy, and Practice*, D. R. Dixon, P. Sturmey, and J. L. Matson, Eds, Cham: Springer Nature Switzerland, 2024, pp. 45–80. doi:10.1007/978-3-031-64499-3_4.
4. American Speech-Language-Hearing Association, 'Social Communication Disorder', American Speech-Language-Hearing Association. Accessed: Jan. 12, 2026. [Online]. Available: <https://www.asha.org/practice-portal/clinical-topics/social-communication-disorder/>
5. C. Hu et al., 'Exploring Speech Pattern Disorders in Autism using Machine Learning', Apr. 10, 2025, Research Square. doi:10.21203/rs.3.rs-6002459/v1.
6. M. Naveed Uddin, 'Cognitive science and artificial intelligence: simulating the human mind and its complexity', *Cognitive Computation and Systems*, vol. 1, no. 4, pp. 113–116, 2019, doi:10.1049/ccs.2019.0022.
7. H. A. Simon, 'Cognitive science: The newest science of the artificial', *Cognitive Science*, vol. 4, no. 1, pp. 33–46, Jan. 1980, doi:10.1016/S0364-0213(81)80003-1.
8. G. M. Biancofiore, D. Di Palma, C. Pomo, F. Narducci, and T. Di Noia, 'Conversational User Interfaces and Agents', in *Human-Centered AI: An Illustrated Scientific Quest*, P. Germanakos, M. Juhasz, A. Kongot, D. Marathe, and D. Sacharidis, Eds, Cham: Springer Nature Switzerland, 2025, pp. 399–438. doi:10.1007/978-3-031-61375-3_4.
9. E. Oliveira, Y. Song, M. Saqr, and S. López-Pernas, 'An Introduction to Large Language Models in Education', in *Advanced Learning Analytics Methods: AI, Precision and Complexity*, M. Saqr and S. López-Pernas, Eds, Cham: Springer Nature Switzerland, 2026, pp. 191–205. doi:10.1007/978-3-031-95365-1_8.
10. A. Fiora and V. Curreri, 'ChatCare: a multimodal conversational interface for people with neurodevelopmental disorders', Apr. 2024, Accessed: Nov. 04, 2025. [Online]. Available: <https://www.politesi.polimi.it/handle/10589/219477?mode=complete>
11. S. Balaji and S. Vigneshwari, 'Comprehensive Assistive Language Model (CALM) for Neurodivergent Child Language Acquisition', in *2025 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI)*, Mar. 2025, pp. 1–7. doi:10.1109/ICDSAAI65575.2025.11011605.
12. X. Zhao et al., 'Beyond one-model-fits-all: A survey of domain specialization for large language models', *arXiv preprint arXiv*, vol. 2305, 2023, Accessed: Nov. 04, 2025. [Online]. Available: <https://www.academia.edu/download/103962038/2305.18703.pdf>

13. G. Balaskas, H. Papadopoulos, D. Pappa, Q. Loisel, and S. Chastin, 'A Framework for Domain-Specific Dataset Creation and Adaptation of Large Language Models', *Computers*, vol. 14, no. 5, p. 172, May 2025, doi:10.3390/computers14050172.
14. W. Lu, R. K. Luu, and M. J. Buehler, 'Fine-tuning large language models for domain adaptation: exploration of training strategies, scaling, model merging and synergistic capabilities', *npj Comput Mater*, vol. 11, no. 1, p. 84, Mar. 2025, doi: 10.1038/s41524-025-01564-y.
15. B. van Aken, 'Exploration and adaptation of large language models for specialized domains', 2023, doi:10.15488/15781.
16. N. Ding et al., 'Parameter-efficient fine-tuning of large-scale pre-trained language models', *Nat Mach Intell*, vol. 5, no. 3, pp. 220–235, Mar. 2023, doi:10.1038/s42256-023-00626-4.
17. Y. Liu, J. Cao, C. Liu, K. Ding, and L. Jin, 'Datasets for large language models: a comprehensive survey', *Artif Intell Rev*, vol. 58, no. 12, p. 403, Oct. 2025, doi: 10.1007/s10462-025-11403-7.
18. P. Röttger, F. Pernisi, B. Vidgen, and D. Hovy, 'SafetyPrompts: A Systematic Review of Open Datasets for Evaluating and Improving Large Language Model Safety', *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 26, pp. 27617–27627, Apr. 2025, doi:10.1609/aaai.v39i26.34975.
19. P. Deshmukh, B. Karmore, M. Ingole, and K. Upreti, 'Designing Emotionally Adaptive Chatbots for Diverse Users: A User-Centered Human-AI Interface Framework', *Turkish Journal of Engineering*, vol. 10, no. 1, pp. 1–12, Oct. 2025, doi: 10.31127/tuje.1715271.
20. ADDA Editorial Team, '10 Signs and Symptoms of ADHD in Adults (And When to Get Help)', ADDA - Attention Deficit Disorder Association. Accessed: Jan. 01, 2026. [Online]. Available: <https://add.org/signs-of-adhd/>
21. CDC, 'Symptoms of ADHD', Attention-Deficit / Hyperactivity Disorder (ADHD). Accessed: Dec. 20, 2025. [Online]. Available: <https://www.cdc.gov/adhd/signs-symptoms/index.html>
22. Cleveland Clinic, 'What Is Cluttering?', Cleveland Clinic. Accessed: Dec. 20, 2025. [Online]. Available: <https://my.clevelandclinic.org/health/diseases/cluttering>
23. J. Chen et al., 'When large language models meet personalization: perspectives of challenges and opportunities', *World Wide Web*, vol. 27, no. 4, p. 42, Jun. 2024, doi:10.1007/s11280-024-01276-1.