Integrating Large Language Models into Therapeutic Education for Children with Dyslexia: a Multimodal Framework

Francesco Piferi[1,*,†], Giovanni Caleffi[1,*,†], Giulia Valcamonica[1,*], Pietro Crovari[1,*] and Franca Garzotto[1]

[1]*Politecnico di Milano*

Abstract

This paper presents ChatCare, a multimodal framework designed to facilitate the development of educational and therapeutic interventions for children with dyslexia, and applicable to enhance interactive educational activities for children with special learning needs. ChatCare is a system that integrates Large Language Models (LLMs) into a chatbot that guides users through personalized linguistic activities to increase the engagement and effectiveness during educational therapy sessions. We describe a modular architecture that integrates emotion detection, adaptive prompt generation and responsive interaction management, enabling natural and context-aware dialogue with young users. The system also features a multimodal interface specifically optimized for children aged 8–11. Designed with a colorful, cartoon-like aesthetic, the interface fosters a playful and engaging environment that aligns with the cognitive and emotional needs of young users. This design not only enhances user experience but also supports the system's core functionalities by accommodating diverse communication preferences and literacy levels through both textual and visual modalities. The paper discusses the design of the architecture, including an overview of the system's rationale and structure, implementation details, and the process leading to the development of prompts to manage with this target population. The main contributions are: (i) a methodology for embedding expert-informed therapeutic or educational prompts into a conversational framework; (ii) an interaction pattern based on a multimodal interface that encourages users to participate in therapeutic activities.

Keywords

Educational Learning, Large Language Models, Dyslexia, Multimodal Framework, Conversational Agent

1. Introduction

Dyslexia is a Neurodevelopmental Disorder (NDD) [1] characterized by reading impairments. Children with dyslexia frequently experience difficulties in phonological processing, working memory, and rapid automatized naming, which can affect other academic domains such as writing and mathematics [2, 3, 4]. Dyslexia requires highly individualized educational and therapeutic strategies [5]. Each child may exhibit unique challenges and learning patterns, which makes traditional approaches more time consuming [5]. To address these challenges, some studies have explored the use of personalized interventions based on interactive digital applications that include continuous feedback and emotional engagement, which have proven beneficial in supporting the learning process [6].
We also witness an increasing interest in using Artificial Intelligence (AI) to enrich educational and therapeutic support for children with special learning needs [7, 8]. In particular, AI-based conversational agents, or chatbots, have shown some potential to increase user engagement in interactive exercises by providing real-time feedback in natural language [9], or to promote specific communication skills [10, 11].

D-SAIL Workshop - Transformative Curriculum Design: Digitalisation, Sustainability, and AI Literacy for 21st Century Learning, July 22, 2025, Palermo, Italy
*Corresponding author.
†These authors contributed equally.
francesco.piferi@polimi.it (F. Piferi); giovanni.caleffi@polimi.it (G. Caleffi); giulia.valcamonica@polimi.it (G. Valcamonica); pietro.crovari@polimi.it (P. Crovari); franca.garzotto@polimi.it (F. Garzotto)
0009-0003-9570-187X (F. Piferi); 0009-0002-9588-0946 (G. Caleffi); 0009-0007-0089-1594 (G. Valcamonica); 0000-0002-6436-4431 (P. Crovari); 0000-0003-4905-7166 (F. Garzotto)

Large Language Models (LLMs) [12] have further expanded the potential for AI-driven educational tools that exploit human-like dialogue and automatically generated personalized content [13]. Research in this field includes studies on user intent and satisfaction with LLM-based assistants [14] and discussions of ethical considerations for chatbots in mental health and learning contexts [15]. While various systems have contributed valuable insights into the use of LLMs for educational and therapeutic support, there remains an opportunity to explore the intersection of three key dimensions: adaptability, emotional awareness and multimodal interaction [16, 5]. In particular, the integration of these features into a cohesive framework tailored for dyslexic children has not been extensively investigated [7]. This paper contributes to fill this gap by presenting the ChatCare framework, an AI-based platform designed to support the development of highly personalized multimodal educational activities. While initially conceived to enhance educational interventions for children with dyslexia, the framework is applicable in principle to a much wider set of educational contexts for children with special learning needs. ChatCare aims to facilitate the creation of a multimodal interaction space for learning that integrates GUI-based interactive exercises for dyslexic children with an advanced conversational agent that assists them during exercise execution. The conversational agent proactively engages children in personalized conversations that take into account both the users' explicit requests (expressed through text or speech) and their "implicit" state, i.e., their emotions and their interactions with the current exercise. To achieve this goal, ChatCare is grounded on a modular architecture, combining LLMs with real-time emotion detection techniques and adaptive context-aware prompt generation. The playful and accessible design of the conversational agent interface further enhances engagement, creating an inclusive environment aligned with the specific cognitive and emotional needs of young users.

Our work brings two primary contributions: (i) we propose a modular system architecture that integrates LLMs, emotion detection, and real-time interaction management into a cohesive educational tool; (ii) we present a multimodal interface optimized for young users, enabling engaging and adaptive educational activities. While we have applied our framework to the treatment of dyslexia, ChatCare can be exploited for a much wider set of educational interventions devoted to children with special learning needs.

The paper is organized as follows: Section 2 reviews related work on educational tools for neurodevelopmental disorders, focusing on dyslexia and Large Language Models. Section 3 details the ChatCare system's architecture, including its modular backend and emotion-aware, multimodal interface. Section 4 and Section 5 cover interaction flows, system adaptability, limitations, and future directions.

2. Related Works

2.1. NDDs Support Systems and Educational Tools

A variety of computer-assisted systems have been developed to support individuals with NDDs, and particularly dyslexia, in both learning and therapy [16, 17, 18, 19]. Early approaches included specialized educational software and assistive technologies focused on specific impairments (e.g., reading software for dyslexic students) [7, 20].

In recent years, conversational agents and social robots have been explored for their ability to engage users. For instance, some systems implemented structured interventions, such as proactive chatbots that enhance engagement by integrating user intent, real-time retrieval, and content generation to deliver personalized experiences [9]. Similarly, chatbot-based interventions have been developed specifically for NDDs, highlighting design considerations for accessibility and family use [11]. In other cases, emotion detection has been introduced to foster more supportive interactions [21, 10]

These systems show how AI can complement traditional therapy by providing additional practice and feedback in a patient's daily environment. While many existing solutions have demonstrated value, they often focus on specific tasks or rely on manually created content [22]. There is still room to expand the variety of interactive activities available, particularly to better support different user needs and therapy stages. Additionally, few platforms currently offer integrated features such as real-time emotional feedback, adaptive difficulty, and multimodal interaction [10, 21].

2.2. Large Language Models in Educational Contexts

Large Language Models like OpenAI's GPT series have brought significant breakthroughs in natural language generation, enabling new possibilities for educational technology [12]. This capability has been quickly used in educational settings: LLM-driven chatbots, for example, may generate practice questions on the fly, clarify topics and mimic conversations [23]. Studies have shown that users find AI assistants helpful for information retrieval and tutoring, though issues like misunderstood intent can occur [14]. In particular, for learners with disabilities, LLMs offer the chance to receive personalized, immediate assistance.

3. ChatCare System Architecture

ChatCare is a modular framework designed to facilitate multimodal, emotionally aware human-computer interaction. Our main purpose was to create a platform in which children with dyslexia, could do useful activities – like reading and writing exercises – supported and tutored by a conversational agent aware of the environment and the context, created ad hoc for the user needs via highly specialized prompt engineering. The motivation behind ChatCare comes from the desire of proposing a more engaging and motivating environment, where children are encouraged to participate in activities that directly support their learning difficulties.
Its architecture consists of a JavaScript-based backend and a React-based frontend, which communicate via HTTP/HTTPS APIs. The backend is responsible for core logic execution, database access through PostgreSQL and integration with the OpenAI API for natural language processing.
ChatCare logic capability is divided into three primary components: the *Chat Manager*, the *Event Manager*, and a set of specialized *Chatbot Modules*. All incoming client requests are first processed by the Chat Manager, which coordinates the interaction between modules and assembles the final query for the LLM, in our case ChatGPT-4o. The Chatbot Modules include:

1. the *Chatbot Rules* module, which defines persistent behavioural guidelines to be followed by the chatbot throughout the conversation. These static prompts enforce conversational consistency and may include constraints based on the user's diagnosis or the specific activity context;
2. the *Emotion Recognition* module, which analyzes both text and audio input from the user to detect underlying emotional states;
3. the *Behavior Prompt Manager*, which tailors the prompt sent to the language model based on the emotional analysis;
4. and the *Clean Input* module, which corrects user input to eliminate syntactic noise and enhance the LLM comprehension [24].

The *Event Manager* was introduced to enable a more proactive and context-aware conversational agent, one that does not rely solely on user-initiated input but can also respond to activity-based events occurring during the session. This component is responsible for monitoring frontend-generated interaction signals, such as task completions, incorrect answers, session initiation, user inactivity, and other interface-level events.
Upon detecting these signals, the *Event Manager* evaluates them according to predefined internal logic and determines whether they warrant additional backend processing. When appropriate, it triggers corresponding actions, such as generating prompts via the *Chat Manager*, which subsequently invokes the LLM to produce chatbot responses.
This architecture allows the system to preserve conversational fluidity and maintain a natural dialogue flow, even when the user is not actively engaging through direct prompts. As a result, the chatbot becomes a more dynamic and human-like agent—capable of offering encouragement, guidance, or re-engagement based on real-time user behavior and activity context.
On the frontend, the *Multimodal Interface Manager* interprets backend responses and distributes them across the graphical user interface (GUI). The GUI consists of two primary views: the *Activity Interface*, where users carry out therapeutic or guided tasks; and the *Conversational Interface*, which features a

chatbot avatar providing contextual support, emotional reinforcement, or clarification when needed. This modular frontend enables an adaptive user experience across both task-driven and dialogue-based interactions.

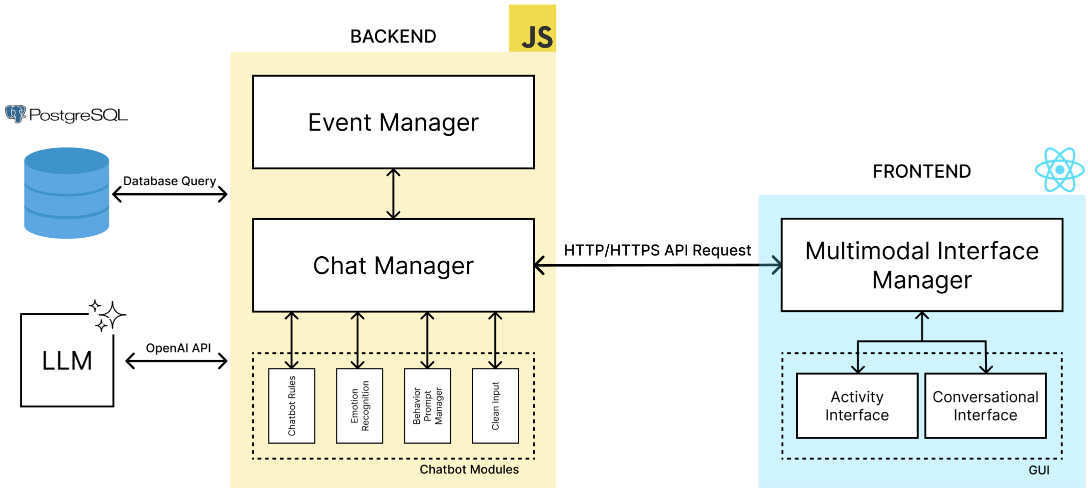

Figure 1: Overview of the ChatCare system architecture, showing backend components for logic, emotion processing, and API integration, and the frontend interface for multimodal interaction.

3.1. Chat Manager

The *Chat Manager* is the core coordination component of the backend architecture. It manages all incoming requests from the frontend and determines the appropriate processing pipeline based on the request type. The system distinguishes between two primary types of input: *user messages* and *interaction events*.

When a user message is received, the Chat Manager initiates a multi-step process involving the Chatbot Modules. It sends the input to the *Clean Input* module for syntactic correction, passes the cleaned message to the *Emotion Detection* module to infer the user's emotional state, and then invokes the *Behavior Prompt Manager* to generate emotion-informed behavioral guidance. It also retrieves persistent conversational rules from the *Chatbot Rules* module. These elements are aggregated into a comprehensive prompt, which is then used to formulate a request for ChatGPT-4o. Session metadata and user interactions are also logged in the PostgreSQL database to maintain state and enable context-aware responses.

In contrast, when the Chat Manager receives an interaction event—such as a task failure or system-triggered cue—it forwards the event to the *Event Manager*. Upon receiving a response that contains a predefined prompt tailored to the user's current context, the Chat Manager directly constructs the API call to the language model using only the prompt provided by the Event Manager, without invoking the Chatbot Modules. This design ensures fast, context-sensitive interventions driven by system logic rather than user input.

3.2. Event Manager

The *Event Manager* processes interaction signals from the frontend, such as repeated task failures, extended user inactivity, or specific interface-triggered events. Each incoming event is evaluated against the *event history*—a log of previously registered frontend events—to determine whether an intervention is required.

For example, if a user incorrectly answers the same exercise item more than a predefined number of times, the Event Manager generates a structured response containing a prompt for the language model. This prompt may instruct the chatbot to clarify the activity, offer guidance, or ask if the user would like help. The Event Manager then forwards this response to the Chat Manager, which uses it directly to construct the query to ChatGPT without invoking the other Chatbot Modules.

We decided to adopt this event-driven, proactive design based on findings that such interventions significantly improve user engagement and prolong conversations in chatbot systems [9].

3.3. Chatbot Modules

The Chatbot Modules are specialized processing units invoked by the Chat Manager to synthesize a rich conversational context before querying the LLM. Each module contributes distinct information essential for delivering personalized and emotionally intelligent interactions:

- **Chatbot Rules**: This module defines static behavioral constraints and role-consistent guidelines, forming a persistent prompt template that the LLM must adhere to throughout the session. These rules may reflect therapeutic goals, user diagnoses, or activity-specific constraints.
- **Emotion Detection**: This module identifies the emotional state of the user based on text and/or audio input. For textual emotion recognition, we use the *Emotion English DistilRoBERTa-base* model [25]. For vocal emotion recognition, we integrate the *Emoty* system [10]. Both models classify emotions using Ekman's six basic categories (anger, disgust, fear, joy, sadness, surprise), along with a neutral class [25].
- **Behavior Prompt Manager**: Once the dominant emotion is identified, this module selects the corresponding behavioral prompt. This prompt provides adaptive conversational strategies tailored to the user's emotional state and is incorporated into the final LLM request improving engagement and empathy from the chatbot [26].
- **Clean Input**: To ensure maximum clarity and relevance in user inputs [24], this module applies advanced grammatical and stylistic corrections using the open-source *LanguageTool* library.

Collectively, these modules allow the system to construct a nuanced, emotionally aware, and behaviorally aligned prompt for the LLM, enhancing the relevance and effectiveness of chatbot responses [25, 26].

3.4. Multimodal Interface

To meet all the functionalities and requirements to create an engaging experience, the frontend of the system was built as a modular, multimodal interface that enables users to engage in both task-based and conversational activities within a unified framework. The user experience is segmented into two synchronized views:

Occupying the left half of the screen, the *Activity Interface* presents therapeutic or guided exercises curated by caregivers. Each activity is semantically encoded and shared with the backend, enabling the chatbot to incorporate detailed knowledge of the user's current task context into its responses. This tight integration fosters continuity and relevance across modalities.

The *Conversational Interface*, located on the right side of the screen, features an animated avatar (Foxy) and a messaging pane supporting both text and audio communication. The chatbot avatar provides emotional reinforcement, task clarification, and conversational support. This component not only responds to user input but also proactively engages based on backend-triggered events or detected emotional states. The avatar's expressions and tone are modulated in alignment with detected emotions, enhancing empathetic resonance.

4. System Behavior and Interaction Flow

To better understand how ChatCare's components work together during operation, we present a flow diagram representing the typical interaction patterns: a user-initiated interaction flow, and an event-driven interaction flow. This diagram traces the steps from the perspective of the user through the backend logic, highlighting how messages and events are handled within the ChatCare architecture.

The interaction flow is handled through two mechanisms: a user-driven dialogue flow, employed for question-and-answer exchanges to deliver immediate, personalized feedback, and an event-driven flow, responsible for managing activity progression and session completion without requiring user prompts.

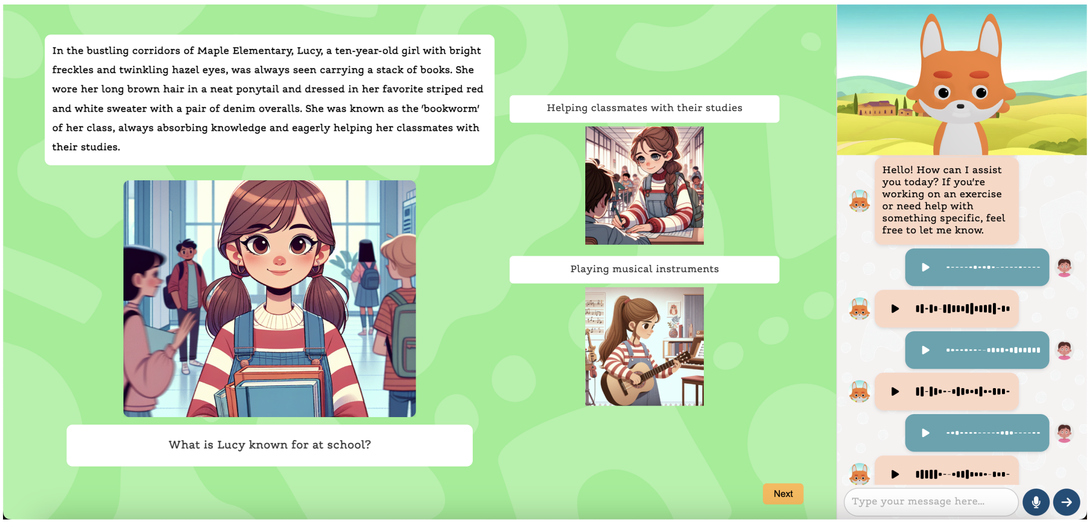

Figure 2: The ChatCare interface combines activity-based exercises with a conversational agent to support user engagement through multimodal interaction.

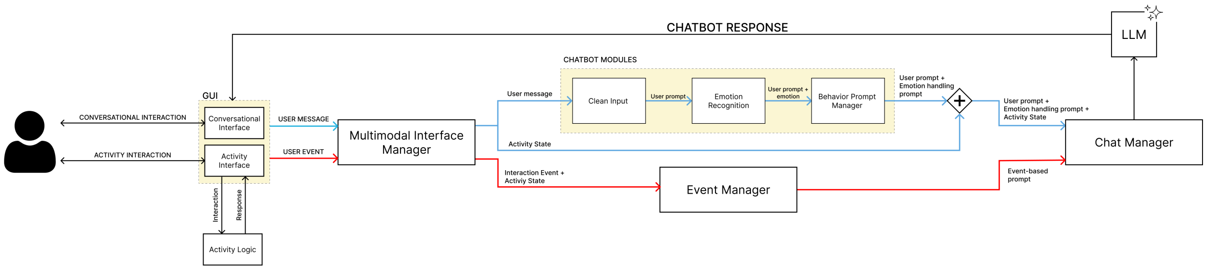

Figure 3: User- and event-driven interaction flow in ChatCare, showing how messages and events are processed to generate LLM responses.

To illustrate the interaction patterns enabled by ChatCare, we simulated a scenario involving a reading session designed for a child with dyslexia. In this session, the caregiver configures a flash spelling activity, where a word is briefly displayed on screen for one second, after which the child must reproduce it with accurate spelling. This scenario demonstrates how ChatCare manages both user-driven dialogue and event-driven transitions.

The session begins when the child logs in. The *Multimodal Interface Manager* sends a Session Start Event together with the current state of the activity to the *Event Manager* for context. The *Event Manager* responds by instructing the *Chat Manager* to trigger an initial greeting. The *Chat Manager* sends a request to ChatGPT, receives the response, and transmits it back to the *Conversational Interface*, which renders the chatbot's message:

> *"Hello! Ready for today's reading practice?"*

In the first flash spelling exercise, the *Activity Interface* displays the word *"kite"* for one second before disappearing. The child types *"kight"* as the response. The *Activity Interface* checks in the *Activity Logic* the correctness of the answer, and since the spelling was incorrect, it highlights the word in red. Internally, this triggers an Incorrect Answer Event, but the *Event Manager* is configured not to initiate chatbot feedback until two incorrect attempts have been made. As a result, the *Chat Manager* at this stage receives the instruction to do nothing.

On the second attempt, the child enters *"kait"*, which is again incorrect. The *Activity Interface* highlights the word in red once more. Upon receiving a second Incorrect Answer Event, the *Event Manager* instructs the *Chat Manager* to issue a prompt. The *Chat Manager* then sends the request to ChatGPT, receives a response, and forwards it to the *Conversational Interface*, which displays the chatbot message:

> *"Almost. The word you are looking for has a long i sound, which is usually written with just an i before the consonant. Try again."*

The child then inputs *"kite"*, which is the correct spelling. The word is highlighted in green by the *Activity Interface*, indicating success. No additional event is triggered, and the activity continues to the next word.
The following word, *"bike"*, is presented for one second and then disappears. Instead of typing a response immediately, the child interacts with the chatbot via voice through the *Conversational Interface*, asking:

> *"Does this word have a long i sound like kite?"*

The *Multimodal Interface Manager* forwards the spoken query to the *Chatbot Modules* to process the input:

- **Clean Input**: Analyzes the content of the message; since it is well-formed, no modifications are made.
- **Chatbot Rules**: Queries the session database to retrieve any configuration-specific prompts needed to contextualize the chatbot's response.
- **Emotion Detection**: Evaluates the audio to identify emotional tone. In this case, the detected emotion is 'Neutral'.
- **Behaviour Prompt Manager**: As the detected emotion is neutral, no additional behavior-based prompts are added to the response.

Once all processing is complete, the *Chat Manager* assembles the final prompt by putting together the output of the *Chatbot Modules* and the activity state, then it makes the LLM call. The resulting response is forwarded and rendered in *Conversational Interface*:

> *"Yes it does! Try writing it!"*

The child enters *"bike"*, which is correctly spelled. The *Activity Interface* highlights the word green and the session continues.
In the end of the assigned activities, the *Multimodal Interface Manager* triggers a Session Complete Event. Following the same internal flow as previous system-driven interactions, the *Chat Manager* sends a final request to ChatGPT and the resulting message is displayed to the user:

> *"You've finished all your exercises for today! You spelled 4 out of 5 words correctly. Excellent work! I'll let your caregiver know how well you did. See you next time!"*

The session transcript is automatically saved and made accessible to the caregiver for post-session review.
This example shows how ChatCare could support educational learning for children with dyslexia through coordinated use of conversational AI, multimodal interfaces, and adaptive feedback.

5. Conclusions and Future Work

This research presented the design and architecture of ChatCare, a multimodal framework created to support educational and therapeutic activities for children with dyslexia. The system integrates a modular backend with a Large Language Model—specifically OpenAI's ChatGPT—alongside emotion detection through both text and voice, and a multimodal user interface that facilitates both structured activities and natural conversation. A simulated use-case scenario was also included to illustrate the system's behavior in practice.
The primary goal of this work is to create a more engaging and responsive environment during educational activities, aiming to better support the needs and participation of children with dyslexia.

To achieve this goal, the paper presented two main contributions: (i) we developed a modular system architecture that integrates Large Language Models with real-time emotion recognition and event-driven interaction management. This architecture was iteratively structured to support adaptive, context-aware dialogue by coordinating multiple backend modules, including specialized prompt engineering, emotion-informed behavioral adjustments, and user interaction monitoring. These components work in concert to create a responsive therapeutic environment aligned with the needs of children with dyslexia. (ii) We designed a multimodal, child-centric interface that balances structured activity delivery with conversational support, employing visual cues tailored to users' cognitive and emotional states. This interface not only facilitates accessibility and engagement but also supports the backend logic by enabling seamless transitions between activity-based and dialogue-based interactions. Through the integration of animated avatars, emotion-sensitive dialogue, and adaptive feedback, the interface aims to encourage a playful yet supportive learning atmosphere.

Despite encouraging architectural advances, ChatCare faces limitations. Integration of third-party tools such as LLMs and the emotion detection models, introduces latency and dependency challenges. The emotion management module's performance, while functional, could benefit from further instruction and validation from experts to meet the children's needs. Scalability remains a concern as the system grows in complexity, particularly when deployed in diverse therapeutic settings with heterogeneous user needs. Ethical considerations surrounding AI-based interventions for vulnerable populations—such as safeguarding privacy and ensuring the appropriateness of feedback—require ongoing attention [15]. While ChatCare appears promising, a real trial on users is necessary to confirm its potential.

Future research is already involving a collaboration with expert linguists to first validate and correct the system prompts. Once done this step, a preliminary explorative study with more than 15 children between 8-11 years old is set to assess the platform's effectiveness and usability in practical therapeutic settings. Moreover, we will aim to enhance the adaptability and generalizability of ChatCare, extending its applicability beyond dyslexia through the integration of additional therapeutic modules for other neurodevelopmental disorders.

Acknowledgments

This work was partially funded by the Ico Falck Foundation under the project "Agenti multimodali intelligenti per interventi valutativi e terapeutici nei bambini con disturbi evolutivi del linguaggio" (Multimodal Intelligent Agents for Assessment and Intervention with Children with Developmental Language Disorders). We thank Denis Delfitto, Maria Verder, and Marta Tagliani for the design and empirical evaluation of the dyslexia exercises.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT-4 and QuillBot for grammar and spelling check. After using these tool(s)/service(s), the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] A. P. Association, Diagnostic and Statistical Manual of Mental Disorders (5th ed.), American Psychiatric Publishing, 2013.

[2] M. McDowell, Specific learning disability, Journal of Paediatrics and Child Health 54 (2018) 1077–1083.

[3] R. L. Peterson, B. F. Pennington, Developmental dyslexia, Lancet 379 (2012) 1997–2007. doi:10.1016/S0140-6736(12)60198-6.

[4] S. E. Shaywitz, B. A. Shaywitz, Dyslexia (specific reading disability), Biological Psychiatry 57 (2003) 1301–1309. doi:10.1016/j.biopsych.2005.01.043.

[5] G. Schulte-Körne, The prevention, diagnosis, and treatment of dyslexia, Deutsches Ärzteblatt International 107 (2010) 718–726. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2967798/. doi:10.3238/arztebl.2010.0718, quiz 727.
[6] E. Kochmar, D. D. Vu, R. Belfer, V. Gupta, I. V. Serban, J. Pineau, Automated personalized feedback improves learning gains in an intelligent tutoring system, 2020. URL: https://arxiv.org/abs/2005.02431. arXiv:2005.02431.
[7] J. R. Yap, T. Aruthanan, M. Chin, Artificial intelligence in dyslexia research and education: A scoping review, IEEE Access 13 (2025) 7123–7134.
[8] L. S. Iyer, T. Chakraborty, K. N. Reddy, K. Jyothish, M. Krishnaswami, Ai-assisted models for dyslexia and dysgraphia: Revolutionizing language learning for children, in: AI-Assisted Special Education for Students With Exceptional Needs, IGI Global, 2023, pp. 186–207.
[9] Z. Niu, Z. Xie, S. Cao, C. Lu, Z. Ye, T. Xu, Z. Liu, Y. Gao, J. Chen, Z. Xu, Y. Wu, Y. Hu, Part: Enhancing proactive social chatbots with personalized real-time retrieval, 2025. URL: https://arxiv.org/abs/2504.20624. arXiv:2504.20624.
[10] F. Catania, N. Di Nardo, F. Garzotto, D. Occhiuto, Emoty: An emotionally sensitive conversational agent for people with neurodevelopmental disorders, 2019. doi:10.24251/HICSS.2019.244.
[11] A. Singla, R. Khanna, M. Kaur, K. Kelm, O. Zaïane, C. S. Rosenfeld, T. A. Bui, F. Majnemer, T. Ogourtsova, Developing a chatbot to support individuals with neurodevelopmental disorders: Tutorial, Journal of Medical Internet Research 26 (2024) e50182.
[12] T. B. Brown, et al., Language models are few-shot learners, 2020. ArXiv preprint arXiv:2005.14165.
[13] J. Chen, Z. Liu, X. Huang, C. Wu, Q. Liu, G. Jiang, Y. Pu, Y. Lei, X. Chen, X. Wang, et al., When large language models meet personalization: Perspectives of challenges and opportunities, World Wide Web 27 (2024) 42.
[14] A. Bodonhelyi, E. Bozkir, S. Yang, E. Kasneci, G. Kasneci, User intent recognition and satisfaction with large language models: A user study with chatgpt, 2024. ArXiv:2402.02136.
[15] S. Coghlan, et al., To chat or not to chat: Ethical issues with using chatbots in mental health, Digital Health 9 (2023) 20552076231183542.
[16] S. Paudel, S. Acharya, A comprehensive review of assistive technologies for children with dyslexia, 2024. URL: https://arxiv.org/abs/2412.13241. arXiv:2412.13241.
[17] G. Morciano, J. M. Alcalde Llergo, A. Zingoni, E. Yeguas Bolívar, J. Taborri, G. Calabrò, Use of recommendation models to provide support to dyslexic students, Expert Systems with Applications 249 (2024) 123738. URL: http://dx.doi.org/10.1016/j.eswa.2024.123738. doi:10.1016/j.eswa.2024.123738.
[18] G. Almgren Bäck, E. Lindeblad, C. Elmqvist, I. Svensson, Dyslexic students' experiences in using assistive technology to support written language skills: a five-year follow-up, Disability and Rehabilitation: Assistive Technology 19 (2024) 1217–1227. doi:10.1080/17483107.2022.2161647.
[19] M. Tagliani, M. Vender, G. Valcamonica, G. Caleffi, F. Garzotto, D. Delfitto, DEA - An Innovative Technological Tool for Personalized Linguistic Training for Italian Children with Developmental Dyslexia, in: S. Rebora, M. Rospocher, S. Bazzaco (Eds.), Diversità, Equità e Inclusione: Sfide e Opportunità per l'Informatica Umanistica nell'Era dell'Intelligenza Artificiale, Proceedings del XIV Convegno Annuale AIUCD2025, Quaderni di Umanistica Digitale, AIUCD, Verona, 2025, p. 663. doi:10.6092/unibo/amsacta/8380.
[20] S. Zhao, S. C. Xiong, B. Pang, X. Tang, P. He, Let ai read first: Enhancing reading abilities for individuals with dyslexia through artificial intelligence, arXiv preprint arXiv:2504.00941 (2025).
[21] A. Brun, R. Liu, A. Shukla, F. Watson, J. Gratch, Exploring emotion-sensitive llm-based conversational ai, 2025. URL: https://arxiv.org/abs/2502.08920. arXiv:2502.08920.
[22] J. Mintz, Y. Huang, Ai in special education: Challenges and opportunities, British Journal of Educational Technology 53 (2022) 1147–1165.
[23] J. C. Farah, S. Ingram, B. Spaenlehauer, F. K. L. Lasne, D. Gillet, Prompting large language models to power educational chatbots, in: H. Xie, C.-L. Lai, W. Chen, G. Xu, E. Popescu (Eds.), Advances in Web-Based Learning – ICWL 2023, volume 14409 of *Lecture Notes in Computer*

Science, Springer, Singapore, 2023. URL: https://doi.org/10.1007/978-981-99-8385-8_14. doi:10.1007/978-981-99-8385-8_14.

[24] B. Wang, C. Wei, Z. Liu, G. Lin, N. F. Chen, Resilience of large language models for noisy instructions, in: Y. Al-Onaizan, M. Bansal, Y.-N. Chen (Eds.), Findings of the Association for Computational Linguistics: EMNLP 2024, Association for Computational Linguistics, Miami, Florida, USA, 2024, pp. 11939–11950. URL: https://aclanthology.org/2024.findings-emnlp.697/. doi:10.18653/v1/2024.findings-emnlp.697.

[25] J. Hartmann, Emotion english distilroberta-base, https://huggingface.co/j-hartmann/emotion-english-distilroberta-base/, 2022.

[26] T. Liu, S. Giorgi, A. Aich, A. Lahnala, B. Curtis, L. Ungar, J. Sedoc, The illusion of empathy: How ai chatbots shape conversation perception, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 39, 2025, pp. 14327–14335.