

and developmental appropriateness—factors that often necessitate simplified instructions, engaging dialogue patterns, and age-appropriate scaffolding to maintain motivation and optimize learning outcomes (Liu et al., 2024c).

To address these challenges, we introduce SingaKids, a dialogic tutor specifically designed to facilitate language learning through picture description tasks. The oral practice enhances children’s language acquisition by stimulating vocabulary development, syntactic complexity, and observational skills, and facilitating contextual language use within meaningful visual contexts - all essential components of early linguistic competence development. To this end, our system integrates four components: (1) dense image captioning to provide rich visual context understanding, (2) multilingual dialogic interaction to support natural conversational flow, and deliver appropriate feedback and guidance, (3) robust speech understanding to process young learners’ verbal responses, and (4) kids-friendly speech generation to improve the student engagement during tutorial sessions. SingaKids operates across four languages relevant to Singapore’s multicultural context: English, Mandarin, Malay, and Tamil, making it accessible to students from diverse linguistic backgrounds.

We further enhanced the system’s performance through multilingual pre-training strategies, task-specific tuning to optimize picture description dialogue flows, and scaffolding optimization to provide appropriate levels of support based on learner responses. This approach allows the system to adapt its interaction patterns to match learners’ proficiency levels and specific linguistic needs. To evaluate the effectiveness of SingaKids, we conducted empirical studies with first and second-grade elementary school students of different language proficiency levels. Our findings demonstrate that the system provides effective dialogic teaching experiences that support language acquisition through natural conversation about visual stimuli. Notably, students at various performance levels showed improvements in descriptive language skills, vocabulary usage, and conversational fluency after engaging with the system.

This work contributes to the growing field of AI-enhanced language education by demonstrating how multimodal, multilingual systems can be successfully deployed to support young learners’ language development. By addressing the challenges of cross-linguistic consistency and age-appropriate

interaction design, SingaKids represents a step forward in creating accessible and effective learning agents for diverse educational contexts.

2 Related Work

Intelligent tutoring systems aim to replicate human tutoring by providing personalized instruction and adaptive feedback to language learners. The advancement of ITSs has marked a significant step forward in education practice (Graesser et al., 2018; Demszky and Hill, 2023; Wang et al., 2023b). These systems provide personalized learning experiences and instant feedback (Chaffar and Frason, 2004; Harley et al., 2015; Grivokostopoulou et al., 2017), tailored to learners’ characteristics and needs (Dzikovska et al., 2014; Grawemeyer et al., 2016; Nihad et al., 2017), and are shown to positively influence students’ engagement in learning and academic performance (Kulik and Fletcher, 2016; Xu et al., 2019).

Dialogue tutor is a particular type of intelligent tutoring system that interacts with students via natural language conversation (Nye et al., 2014; Ruan et al., 2019). In STEM domains, conversational ITSs can facilitate university students in problem-solving by providing real-time feedback and hints in text formats (Nye et al., 2023; Paladines and Ramirez, 2020; Arnau-González et al., 2023). However, prior work has widely relied on rule-based systems with human-crafted domain knowledge (Nye et al., 2014; Graesser et al., 2018), or data-driven approaches that require a certain amount of human annotation for supervised learning (MacLellan and Koedinger, 2022). Recently, LLMs show strong potential to build dialogue tutors with less data supervision and higher coherence (Afzal et al., 2019; Demszky and Hill, 2023; Macina et al., 2023b), and they can be further improved by integrating LLMs with pedagogical and learning science principles (Stasaski et al., 2020; Sonkar et al., 2023; Macina et al., 2023a).

3 SingaKids System Architecture

In a picture description session, teachers first present an image and ask students to observe it carefully. They pose open-ended questions like “*What do you see in this picture?*” to stimulate observation, then guide students beyond basic object identification to describe qualities using adjectives and adverbs, enhancing vocabulary, organization, and fluency. The activity concludes with introduc-

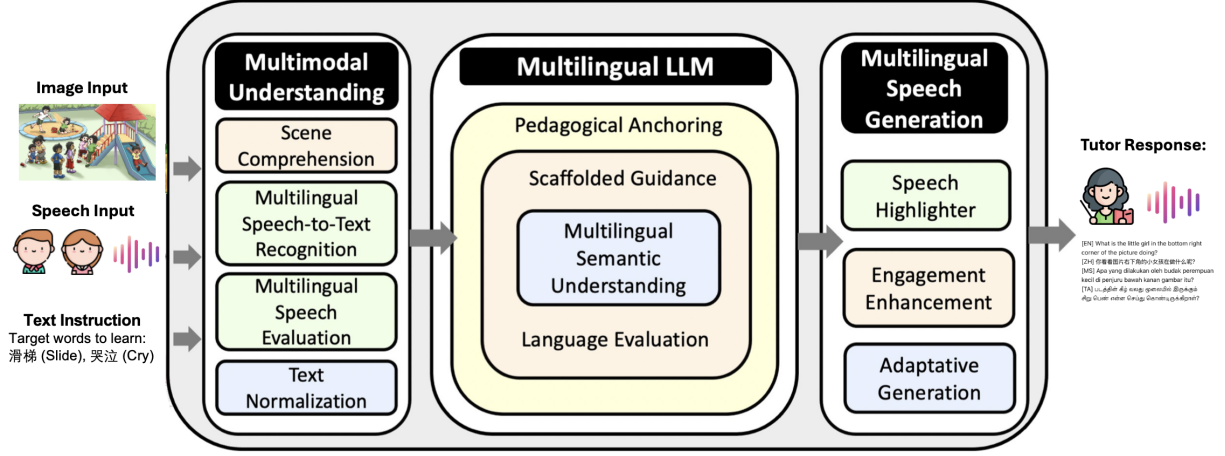


Figure 2: Overview of the conversational tutor architecture for language learning via the image description task.

ing new vocabulary and encouraging students to create stories about the image, developing creativity and narrative skills.

Drawing from real-world teaching sessions, the overall architecture of our conversational tutoring system is illustrated in Figure 2. Interactions begin with the Multimodal Understanding: the scene comprehension extracts the keywords, objects, and events in the given picture (Johnson et al., 2016); the multilingual speech recognition converts the student’s spoken response into text; speech evaluation component is to assess the student’s oral language proficiency (Wong et al., 2022).

The Multilingual LLM represents the core of educational interaction: multilingual semantic understanding interprets the student’s response in context; language evaluation assesses the linguistic accuracy and completeness of their description; scaffolded guidance determines the appropriate level of support needed; This component effectively analyzes the student’s current understanding and formulates an appropriate teaching strategy; pedagogical anchoring establishes high-level educational objectives such as word understanding or sentence construction. Moreover, for elementary grade 1 and 2, we evaluate students’ skills in making sentences to describe the activities in the image, focusing on their vocabulary usage. The language evaluation can be adapted for higher grade levels, by measuring grammatical correctness and coherent narratives (Genishi and Dyson, 2015).

The system’s response is formed in both text and audio outputs. The Multilingual Speech Generation converts text utterance into natural and engaging speech to maintain student motivation (Kim et al., 2021); In addition, beyond simple text-to-speech

synthesis, we incorporate a highlight component which can emphasize important keywords or pronunciation errors (Zhang et al., 2021).

Throughout this pipeline, the system maintains an educational dialogic flow, asking guiding questions, providing hints, offering corrections, and acknowledging progress as needed. If a student struggles with specific vocabulary when describing the image (for example, using general terms like “playing” instead of specific verbs like “swimming” or “climbing”), the system will scaffold their learning through targeted questions and gradually decreasing support until they can independently produce the target words (Liu et al., 2024c).

4 Module Optimization

Young students at the early elementary stage with limited language proficiency raise unique requirements for human-AI interaction. Enhancing the multilingual capability of the core components can improve communication efficacy as well as handling mixed language usage or intra-sentential code-switching. This is particularly important in environments where learners may express themselves in multiple languages they are exposed to at home or school. Additionally, scaffolding kids requires simplified instructions, and consistent engagement through positive feedback and social-emotional support. While maintaining reasonable performance in English and Mandarin, we specifically focus on improving Malay and Tamil to better serve Singapore’s multilingual student population.¹

¹We used the Huggingface codebase for model training & evaluation (<https://github.com/huggingface/transformers>). All experiments were conducted on Nvidia A100 40/80GB GPUs.




Figure 3: Dense image captioning with contextualization (caption of each event aligns with large narrative) and fine-grained understanding (detailed description of objects, characters and activities).

4.1 Fine-grained Image Description

For the picture-guided conversation flow, we propose a dense image captioning module for visual storytelling. The goals are to identify the key events of interest in the image as well as generate rich captions for each key event of interest. Referring to the example in Figure 3, the caption for each event shall be aligned with the larger narrative of the image (better contextualization), and include detailed description of the objects, characters, and activities (fine-grained understanding). State-of-the-art multi-modal LLMs (MLLMs), especially smaller size models, generally struggle with dense image content. The MLLMs often generate general and broad descriptions of the image, and are limited in deeper analysis of the visual details. Moreover, hallucinations occur due to complex or ambiguous image content. Hence, we adopt a two-stage approach – event bounding box proposal and caption generation. For event bounding box proposal, we leverage robust person and object detection (Liu et al., 2024a), human segmentation (Kirillov et al., 2023), coupled with depth estimation (Bhat et al., 2023), for probabilistic reasoning on the neural detections. For caption generation, we use chain-of-thought prompting on a MLLM, InternVL2.5 (Chen et al., 2024), to incorporate global context understanding into the individual event captions.




Figure 4: Malay ASR evaluation results.




Figure 5: Tamil ASR evaluation results.

We achieved a 75% sentence-level accuracy in our image testbed, which provides reasonable content for the conversational process.

4.2 Improving Multilingual ASR

To enhance the multilingual ASR capabilities of our system, we selected Whisper-large-V3 (Radford et al., 2022) as the base model and fine-tuned it with additional Malay and Tamil speech data. Preliminary analysis revealed significant performance gaps when processing lower-resource languages (e.g., Malay, Tamil), and in children’s voice transcribing (Attia et al., 2024). We addressed this limitation by gathering a local dataset comprising 2,800-hour Tamil recordings and 1,000-hour Malay recordings from more than 1,000 native speakers from different age groups and linguistic contexts.

As shown in Figure 4 and Figure 5, we compare Whisper-large-V3 with our fine-tuned model on Malay and Tamil data. Evaluating on Malay data, we achieved a lower WER from 40.5% to 28.4% on conversational speech and from 20.3% to 5.1% on children speech (Zhang et al., 2021). For Tamil, we achieved lowers WER from 10.3% to 7.1% on Bloom Speech Tamil (Leong et al., 2022) and from 13.7% to 7.9% on a children speech data (Zhang et al., 2021). We obtained consistent improvements at all test sets, particularly in children’s voice.




Figure 6: Comparison between the base model and our model of translation capability.




Figure 7: Comparison between the base model and our model of multilingual instruction following.

4.3 Improving Dialogue LLM

We improve the dialogue component built on a text LLM from the following two aspects:

4.3.1 Multilingual Capability

LLMs often show downgraded performance in low-resource languages, and this problem becomes more prominent on smaller models. In this work, we selected Qwen1.5-4B (Bai et al., 2023) as the base model for a balance of performance and cost-efficiency.² Our multilingual optimization follows a two-stage process: First, we conducted continue pre-training on 14B tokens of mixed data with four languages (English, Mandarin, Malay, Tamil) (Penedo et al., 2024). We set a balanced sampling rate to elevate the multilingual modeling of Malay and Tamil, and English and Mandarin data play a role to retain the fundamental language capabilities. Second, we enhanced the model’s multilingual instruction following by multi-task learning (Teknium, 2023) and cross-lingual alignment (Muennighoff et al., 2023; Lin et al., 2025), including multilingual role-play corpora generated through simulating diverse conversation scenarios (Sun et al., 2024; Liu et al., 2024b). To further

²We tested a set of Qwen1.5 models from 1.8B to 14B, and observed that model size is strongly correlated with multilingual capabilities, especially for languages with lower resources such as Malay and Tamil.




Figure 8: Comparison between the base dialogue model and pedagogical-enhanced model through the LLM-as-a-judge evaluation.

strengthen cross-lingual capabilities, we did a hybrid training approach that combines translation and cross-lingual problem-solving tasks (Muennighoff et al., 2022; Liu et al., 2022). This enables the language model of better semantic fusion across languages. Experimental results shown in Figure 6 and Figure 7 show improvement on multilingual translation and instruction following.

4.3.2 Scaffolding-guided Augmentation

We improved the dialogue model’s pedagogical effectiveness by training with scaffolding instructions and personality-aware student simulation (de Oliveira et al., 2023; Sonkar et al., 2023; Liu et al., 2024c,d). Our scaffolding framework is formulated upon the dialogic teaching theory (Alexander, 2006), where the tutor encourages exchange of ideas using follow-up questions, clues, elaborations, or recaps. We conducted a theory-inspired practice by sampling synthetic dialogue samples from a stronger teacher LLM (GPT-4 (Achiam et al., 2023b)) to guide the smaller LLMs, which is capable of providing scaffolded interactions based on learners response. Moreover, in dialogic teaching, recognizing and adapting to individual characteristics can significantly enhance student engagement and learning efficiency. We built a taxonomy of student personality profiles based on established traditional psychology frameworks (Costa and McCrae, 1999), and integrated both cognitive and noncognitive aspects into LLM-based personality-aware student simulation (Liu et al., 2024d). This augmentation enabled the dialogue model to dynamically adjust its pedagogical approach, providing encouragement for students exhibiting low confidence or being distracted from the on-going session, as shown in Figure 8.

Moreover, we observed that the scaffolding-guided training improves the dialogue model’s ro-




Figure 9: Subjective evaluation results for mean opinion score (MOS) and objective evaluation results for speech intelligibility with 95% confidence intervals for the synthesized Malay and Tamil speech samples by the proposed educational multilingual TTS.

business regarding inappropriate language and random user inputs. By incorporating dialogic teaching principles and personality-aware student simulation, our system maintains focus on educational objectives and avoids generating harmful or off-topic responses. For instance, when faced with unexpected user behaviors, the model usually prompts the students back to the image description task (i.e., adopting the scaffolding type “instruction”).

4.4 Improving Multilingual TTS

For engaging speech generation, we utilized VITS (Kim et al., 2021), a non-autoregressive framework that achieves a balance between speech quality and computational efficiency. In particular, for low-resource scenarios Malay and Tamil, we collected recordings from adult teachers and children for modeling appropriate prosodic patterns and speech rhythms. The Malay training data includes 22 hours of adult speech from 1 speaker and 9 hours of child speech from 97 speakers, while the Tamil training data comprises 63 hours of adult speech from 1 speaker and 1.5 hours of child speech from 52 speakers. Speaker embeddings are in a one-hot input format, followed by embedding layer, enabling multi-speaker generation. This approach addresses the issue of voice naturalness in educational contexts, as our preliminary testing revealed that students engage more effectively with systems that generate age-appropriate speech.

As shown in Figure 9, both objective and subjective evaluations are conducted to assess the multilingual TTS performance. Subjective evaluation is conducted using Mean Opinion Score (MOS) ratings, where listeners assess the overall speech




Figure 10: Dialogue analysis on scaffolding types of high-performing and low-performing students.

quality and naturalness of the synthesized speech on a 1-5 scoring. It includes 15 Malay and 15 Tamil child speech samples, rated by 20 native Malay and Tamil listeners respectively. The MOS results indicate that our TTS models achieved a reasonable performance (with an average score exceeding 3.50), showing the effectiveness of multilingual adaptation on naturalness and overall quality. For objective evaluations, we measure speech intelligibility using the widely adopted character error rate (CER). Specifically, we used pretrained Malay and Tamil ASR models to transcribe the Malay and Tamil TTS generation, and computed the CER to quantify speech intelligibility. The results (see Figure 9) demonstrate that our TTS models achieve high speech intelligibility, with recognition accuracy exceeding 90% for both Malay and Tamil.

5 Student Practice and Discussion

We conducted a user study with 35 elementary school students (grade 1-2) to evaluate SingaKids’ effectiveness in real-world educational settings (IRB number: IRB-2024-218). Participants represented diverse language proficiency levels, and they were using the system under the consent and guidance from their parents. Following previous work (Liu et al., 2024c), we conducted a utterance-level analysis of the 7 scaffolding types. As shown in Figure 10, significant differences are in some scaffolding types. High-performing students receive more feeding back (69% vs. 43%) and explanations (21% vs. 9%), where students were encouraged toward deeper understanding. Low-performing students received more hints (12% vs. 5%) and social-emotional support (31% vs. 17%); the system provides clues, support building confidence when learners struggle.

Moreover, there are observations from our preliminary study: (1) In some cases, noisy environments and children’s speech led to more ASR errors, affecting the communication quality. Noise-robust speech recognition and speaker recognition and diarization can help mitigate these issues; (2) Even with dynamic scaffolding and social-emotional support, some students exited sessions when facing persistent difficulties. The scaffolding type “*Modeling*” needs to be triggered to prevent frustration; (3) For lower elementary grades, parent guidance is necessary, as they can provide assistance and additional support; (4) When there are many objects and activities in the picture, kids sometimes become distracted or have difficulty pinpointing the focus area, and adding visual highlighting (e.g., bounding boxes) helps improve focus and comprehension. These findings underline the importance of modeling kids-specific learning preferences, to create a more inclusive and effective language learning experience.

6 Conclusion

In this work, we presented SingaKids, a multilingual multimodal dialogic tutor designed to enhance elementary language acquisition through picture description tasks. By integrating dense image captioning, multilingual interaction, speech understanding, and engaging speech generation across four languages, our system creates an interactive learning environment that adapts to diverse linguistic contexts. Considering the speech and language proficiency and learning objectives of elementary students, we further improved the system on task-specific optimization and age-appropriate pedagogical alignment. Preliminary empirical studies with elementary school students demonstrated SingaKids’ effectiveness in providing self-adaptive guidance through dynamic scaffolding and social-emotional support. Our work provides both technical and educational insights to build general agents in broader educational contexts.

Limitations

We are aware that it remains an open problem to mitigate hallucinations and biases in large language models, which may cause communication issues in human-machine interaction and computer-assisted education. Of course, current models and laboratory experiments are always limited in this or similar ways. We do not foresee any unethical uses

of our proposed methods or their underlying tools, but hope that it will contribute to reducing incorrect system outputs.

Ethics and Impact Statement

We acknowledge that all of the co-authors of this work are aware of the provided ACL Code of Ethics and honor the code of conduct. In our experiments, models are applied under proper license. All data used in this work are only for academic research purposes and should not be used outside of academic research contexts. Our proposed methodology in general does not create a direct societal consequence and are intended to be used to improve the performance, robustness, and safety of the intelligent tutoring systems.

Acknowledgments

This research is supported by the AI4EDU Programme in the Institute for Infocomm Research (I²R), Agency for Science, Technology and Research (A*STAR), and the National Research Foundation, Singapore under its AISG Programme (AISG2-GC-2022-005). We thank the anonymous reviewers for their precious feedback to help improve and extend this piece of work. We acknowledge valuable support and assistance from Siti Umairah Md Salleh, Siti Maryam Binte Ahmad Subaidi, Nabilah Binte Md Johan, Amudha Narayanan, and Anitha Veeramani at the Institute for Infocomm Research (I²R), and valuable contribution in research discussion and study coordination from Chong Han, Audi Arwani Binte Azlan, and Sumi Baby Thomas at the National Institute of Education (NIE), Singapore.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023a. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023b. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Shazia Afzal, Tejas Dhamecha, Nirmal Mukhi, Renuka Sindhgatta, Smit Marvaniya, Matthew Ventura, and Jessica Yarbrow. 2019. [Development and deployment](#)

- of a large-scale dialog-based intelligent tutoring system. In *Proceedings of the NAACL 2019*, pages 114–121, Minneapolis, Minnesota. Association for Computational Linguistics.
- Robin Alexander. 2006. *Education as Dialogue: Moral and Pedagogical Choices for a Runaway World*. Hong Kong Institute of Education.
- Pablo Arnau-González, Miguel Arevalillo-Herráez, Romina Albornoz-De Luise, and David Arnau. 2023. A methodological approach to enable natural language interaction in an intelligent tutoring system. *Computer Speech & Language*, 81:101516.
- Ahmed Adel Attia, Jing Liu, Wei Ai, Dorottya Demszky, and Carol Espy-Wilson. 2024. Kid-whisper: Towards bridging the performance gap in automatic speech recognition for children vs. adults. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, pages 74–80.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. 2023. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*.
- Soumaya Chaffar and Claude Frasson. 2004. Inducing optimal emotional state for learning in intelligent tutoring systems. In *International Conference on Intelligent Tutoring Systems*, pages 45–54. Springer.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24185–24198.
- Paul T Costa and Robert R McCrae. 1999. A five-factor theory of personality. *The five-factor model of personality: Theoretical perspectives*, 2:51–87.
- Luciana C de Oliveira, Loren Jones, and Sharon L Smith. 2023. Interactional scaffolding in a first-grade classroom through the teaching–learning cycle. *International Journal of Bilingual Education and Bilingualism*, 26(3):270–288.
- Dorottya Demszky and Heather Hill. 2023. The ncte transcripts: A dataset of elementary math classroom transcripts. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 528–538.
- Myroslava Dzikovska, Natalie Steinhauser, Elaine Farrow, Johanna Moore, and Gwendolyn Campbell. 2014. BEETLE II: Deep Natural Language Understanding and Automatic Feedback Generation for Intelligent Tutoring in Basic Electricity and Electronics. *International Journal of Artificial Intelligence in Education*, 24(3):284–332.
- Celia Genishi and Anne Haas Dyson. 2015. *Children, language, and literacy: Diverse learners in diverse times*. Teachers College Press.
- Arthur C Graesser, Xiangen Hu, and Robert Sottolare. 2018. Intelligent tutoring systems. In *International handbook of the learning sciences*, pages 246–255. Routledge.
- Beate Grawemeyer, Manolis Mavrikis, Wayne Holmes, Sergio Gutierrez-Santos, Michael Wiedmann, and Nikol Rummel. 2016. Affecting off-task behaviour: how affect-aware feedback can improve student learning. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge, LAK '16*, page 104–113, New York, NY, USA. Association for Computing Machinery.
- Foteini Grivokostopoulou, Isidoros Perikos, and Ioannis Hatzilygeroudis. 2017. An Educational System for Learning Search Algorithms and Automatically Assessing Student Performance. *International Journal of Artificial Intelligence in Education*, 27(1):207–240.
- Jason M. Harley, François Bouchet, M. Sazzad Hussain, Roger Azevedo, and Rafael Calvo. 2015. A multi-componential analysis of emotions during complex learning with an intelligent multi-agent system. *Computers in Human Behavior*, 48:615–625.
- Hyangeun Ji, Insook Han, and Yujung Ko. 2023. A systematic review of conversational ai in language education: Focusing on the collaboration with human teachers. *Journal of Research on Technology in Education*, 55(1):48–63.
- Justin Johnson, Andrej Karpathy, and Li Fei-Fei. 2016. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Jaehyeon Kim, Jungil Kong, and Juhee Son. 2021. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *International Conference on Machine Learning*, pages 5530–5540. PMLR.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo,

- et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026.
- James A. Kulik and J. D. Fletcher. 2016. [Effectiveness of Intelligent Tutoring Systems: A Meta-Analytic Review](#). *Review of Educational Research*, 86(1):42–78.
- Colin Leong, Joshua Nemecek, Jacob Mansdorfer, Anna Filighera, Abraham Owodunni, and Daniel White-nack. 2022. Bloom library: Multimodal datasets in 300+ languages for a variety of downstream tasks. *arXiv preprint arXiv:2210.14712*.
- Geyu Lin, Bin Wang, Zhengyuan Liu, and Nancy Chen. 2025. Crossin: An efficient instruction tuning approach for cross-lingual knowledge alignment. In *Proceedings of the Second Workshop on Scaling Up Multilingual & Multi-Cultural Evaluation*, pages 12–23.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. 2024a. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, pages 38–55. Springer.
- Zhengyuan Liu, Shikang Ni, Aiti Aw, and Nancy Chen. 2022. Singlish message paraphrasing: A joint task of creole translation and text normalization. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3924–3936.
- Zhengyuan Liu, Stella Xin Yin, and Nancy Chen. 2024b. [Optimizing code-switching in conversational tutoring systems: A pedagogical framework and evaluation](#). In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 500–515, Kyoto, Japan. Association for Computational Linguistics.
- Zhengyuan Liu, Stella Xin Yin, Carolyn Lee, and Nancy F Chen. 2024c. Scaffolding language learning via multi-modal tutoring systems with pedagogical instructions. In *2024 IEEE Conference on Artificial Intelligence (CAI)*, pages 1258–1265. IEEE.
- Zhengyuan Liu, Stella Xin Yin, Geyu Lin, and Nancy F. Chen. 2024d. [Personality-aware student simulation for conversational intelligent tutoring systems](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 626–642, Miami, Florida, USA. Association for Computational Linguistics.
- Jakub Macina, Nico Daheim, Sankalan Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023a. Mathdial: A dialogue tutoring dataset with rich pedagogical properties grounded in math reasoning problems. In *Findings of EMNLP 2023*, pages 5602–5621.
- Jakub Macina, Nico Daheim, Lingzhi Wang, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023b. Opportunities and challenges in neural dialog tutoring. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2357–2372.
- Christopher J MacLellan and Kenneth R Koedinger. 2022. Domain-general tutor authoring with apprentice learner models. *International Journal of Artificial Intelligence in Education*, 32(1):76–117.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, et al. 2023. Crosslingual generalization through multitask finetuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, et al. 2022. Crosslingual generalization through multitask finetuning. *arXiv preprint arXiv:2211.01786*.
- Elghouch Nihad, En-naimi El Mokhtar, and Yassine Zaoui Seghroucheni. 2017. [Analysing the outcome of a learning process conducted within the system als_corr\(lp\)](#). *International Journal of Emerging Technologies in Learning (iJET)*, 12(03):pp. 43–56.
- B Nye, Dillon Mee, and Mark G Core. 2023. Generative large language models for dialog-based tutoring: An early consideration of opportunities and concerns. In *AIED Workshops*.
- Benjamin D Nye, Arthur C Graesser, and Xiangen Hu. 2014. Autotutor and family: A review of 17 years of natural language tutoring. *International Journal of Artificial Intelligence in Education*, 24:427–469.
- José Paladines and Jaime Ramirez. 2020. A systematic literature review of intelligent tutoring systems with dialogue in natural language. *IEEE Access*, 8:164246–164267.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben al-lal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. [The fineweb datasets: Decanting the web for the finest text data at scale](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Silvia Pokrivčáková. 2019. Preparing teachers for the application of ai-powered technologies in foreign language education. *Journal of language and cultural education*.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#).

- Sherry Ruan, Liwei Jiang, Justin Xu, Bryce Joe-Kun Tham, Zhengneng Qiu, Yeshuang Zhu, Elizabeth L Murnane, Emma Brunskill, and James A Landay. 2019. Quizbot: A dialogue-based adaptive learning system for factual knowledge. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–13.
- Shashank Sonkar, Naiming Liu, Debshila Mallick, and Richard Baraniuk. 2023. Class: A design framework for building intelligent tutoring systems based on learning science principles. In *Findings of EMNLP 2023*, pages 1941–1961.
- Katherine Stasaski, Kimberly Kao, and Marti A Hearst. 2020. Cima: A large open access dialogue dataset for tutoring. In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–64.
- Yuchong Sun, Che Liu, Kun Zhou, Jinwen Huang, Ruihua Song, Wayne Xin Zhao, Fuzheng Zhang, Di Zhang, and Kun Gai. 2024. Parrot: Enhancing multi-turn instruction following for large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9729–9750.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Teknum. 2023. [Openhermes 2.5: An open dataset of synthetic data for generalist llm assistants](#).
- Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, Ai Ti Aw, and Nancy F Chen. 2023a. SeaEval for multilingual foundation models: From cross-lingual alignment to cultural reasoning. *arXiv preprint arXiv:2309.04766*.
- Rose E Wang, Qingyang Zhang, Carly Robinson, Susanna Loeb, and Dorottya Demszky. 2023b. Step-by-step remediation of students’ mathematical mistakes. *arXiv preprint arXiv:2310.10648*.
- Jeremy Heng Meng Wong, Huayun Zhang, and Nancy F Chen. 2022. Variations of multi-task learning for spoken language assessment. In *Interspeech*, pages 4456–4460.
- Zhihong Xu, Kausalai Wijekumar, Gilbert Ramirez, Xueyan Hu, and Robin Irey. 2019. [The effectiveness of intelligent tutoring systems on K-12 students’ reading comprehension: A meta-analysis](#). *British Journal of Educational Technology*, 50(6):3119–3137.
- L. Yan, L. Sha, L. Zhao, Y. Li, R. Martinez-Maldonado, G. Chen, X. Li, Y. Jin, and D. Gašević. 2024. [Practical and ethical challenges of large language models in education: A systematic scoping review](#). *British Journal of Educational Technology*, 55(1):90–112.
- Xuesong Zhai, Xiaoyan Chu, Ching Sing Chai, Morris Siu Yung Jong, Andreja Istenic, Michael Spector, Jia-Bao Liu, Jing Yuan, and Yan Li. 2021. A review of artificial intelligence (ai) in education from 2010 to 2020. *Complexity*, 2021(1):8812542.
- Huayun Zhang, Ke Shi, and Nancy F Chen. 2021. Multilingual speech evaluation: Case studies on english, malay and tamil. In *Proc. Interspeech 2021*, pages 4443–4447.
- Ke Zhang and Ayse Begum Aslan. 2021. Ai technologies for education: Recent research & future directions. *Computers and education: Artificial intelligence*, 2:100025.