Undergraduate Thesis Prospectus

Adapting Embodied Agents for Creativity: Benchmarking in Virtual Worlds

(technical research project in Computer Science)

The Place of AI in Early Childhood Education: A Divisive Debate

(sociotechnical research project)

by

Yifan Zhou

November 8, 2024

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Yifan Zhou

Technical advisor: Professor Yen-Ling Kuo, Department of Computer Science

STS advisor: Professor Peter Norton, Department of Engineering and Society

Prospectus

General Research Problem

Can AI systems exhibit creativity in virtual world environments?

The year 2022 marked a defining shift in generative AI, transforming it from experimental technology into a mainstream force. Early in the year, ChatGPT's release introduced large-scale conversational AI capable of sophisticated language tasks. This model demonstrated impressive capacity for text generation, quickly becoming a benchmark for generative AI's creative potential in dialogue, content creation, and storytelling. ChatGPT's ability to engage in language-driven creativity underscored AI's potential across writing, poetry, and technical domains.

Soon after, DALL-E emerged as a breakthrough in visual generative AI. This model achieved remarkable precision in generating images from text, setting a new standard for AI-driven artistry. DALL-E's release broadened AI's creative impact into visual arts, establishing a foundation for future progress in image generation and inspiring widespread interest in AI-fueled artistic creation.

As 2022 unfolded, generative AI advanced into multimodal realms, incorporating text, audio, and visual inputs to create more immersive, multi-sensory experiences. This evolution suggested AI could soon generate holistic virtual environments that mimic human creativity across media, signaling a leap from single-medium tasks to integrated creative experiences.

These advances sparked societal reflection on AI's role in creative industries. Some feared AI's capacity to disrupt traditional arts by competing with human creators, while others saw it as an enhancer of human creativity, accelerating and amplifying innovative work. This debate

underscored the need to examine AI's place in the creative landscape carefully, balancing its potential benefits with ethical considerations.

Research began exploring the extent to which AI can achieve authentic creativity, particularly within complex virtual worlds that demand integrated sensory engagement. Studies focused on developing AI agents capable of supporting human creativity, imagining scenarios where AI could enhance educational and social experiences. Potential applications include AI companions that foster creativity in early childhood through interactive play or serve as supportive aides for the elderly. The idea of AI as a lifelong creative partner points to a future where AI could play an integral role in learning, creativity, and personal growth across all ages.

Technical Research Project

How can open-ended embodied agents be adapted to demonstrate and benchmark their creative problem-solving abilities in complex virtual environments such as Minecraft?

Artificial intelligence (AI), one of today's most widely discussed topics, intersects closely with cognition psychology—the branch of psychology focused on understanding internal mental processes such as perception, memory, reasoning, and problem-solving. Cognition psychology explores how humans acquire, process, and store information, providing insights into behavior and decision-making. This overlap between AI and cognition psychology enriches both fields, as insights into human cognition inform AI design and reveal potential for adaptive, human-like decision-making in AI.

Inspired by creative projects like Mindojo (Fan et al., 2022) and Mindcraft (Kolbytn, n.d.), This technical research project leverages this interdisciplinary perspective to investigate the creative potential of AI agents within open-world, gamified environments like Minecraft. By

examining how large language model (LLM) AI agents—such as ChatGPT, Gemini, Claude, and Llama—tackle complex, open-ended tasks, the study aims to benchmark and quantify their creativity. Minecraft's flexible, interactive world offers an ideal platform for assessing AI creativity, as it allows agents to explore solutions freely, simulating conditions that encourage innovation and adaptive problem-solving. The goal is to benchmark each agent's performance and explore how AI creatively engages with tasks that reflect human cognitive processes, ultimately assessing the potential for AI agents to collaborate with humans on tasks in more efficient and innovative ways.

Drawing inspiration from established research, particularly the Torrance Tests of Creative Thinking (TTCT) (Torrance, n.d.) and the Consensual Assessment Technique (CAT) (Cseh & Jeffries, 2019), this project adopts a similar benchmarking approach as used in "Judging the creative prowess of AI" (Chakraborty & Masud, 2023; Torrance, n.d.). Each AI agent is given a standardized set of prompts detailing tasks such as building a house with specific features, decorating a structure, designing a garden, and navigating a maze. These tasks are crafted to challenge the agent's problem-solving and creative thinking abilities, allowing for comparative analysis across different AI models. Each AI agent follows an identical prompt structure to ensure consistency in task requirements. No roleplay for agents is required; the agents focus solely on completing the tasks to showcase their creativity and adaptability. The AI agents perform these tasks autonomously within the Minecraft environment, and their actions, completion times, and outputs are recorded for analysis.

The evaluation focuses on three main aspects of agent performance: Task Completion Time, Output Quality, and Material Diversity. Task Completion Time measures how long each agent takes to complete each task, while Output Quality assesses the structural integrity and logic of each output, ensuring, for example, that a "house" includes recognizable elements like doors,

windows, and a roof. Material Diversity examines the variety in building materials, colors, and textures used by the agents to determine if they demonstrate creativity through material selection and aesthetic appeal. For maze navigation tasks, the agent's chosen path is analyzed to evaluate its efficiency and problem-solving strategy, including the ability to bypass obstacles and find shorter routes.

Once tasks are completed, each output is evaluated using creativity assessment criteria based on psychological benchmarks like the TTCT and CAT. This evaluation includes metrics such as originality, complexity, usefulness, and surprise. Each agent's performance receives a quantified "creative score," ranging from 0 to 10, which indicates their level of creativity across tasks. The resulting creative scores for each agent are then compared to determine relative creativity levels, enabling a systematic evaluation of the agents' strengths and limitations in problem-solving and creative adaptability. This project aims to establish a comparative framework for evaluating AI creativity in virtual environments, providing key insights into AI's creative potential for both simulated and real-world applications.

Sociotechnical Research Project

How are educators, parents, and tech companies competing to determine AI's proper place in early childhood education?

Educators, parents, and tech companies are competing to determine AI's place in early childhood education, each with unique priorities and concerns. At a recent Stanford event, Sam Altman noted AI's appeal to younger users, observing that while AI excels in fields like chess, teenagers increasingly prefer interacting with AI therapists over human ones. This shift highlights AI's growing influence and interest in its role in early education. For participants, the potential benefits and risks of AI in early childhood education bring complex issues, from privacy and developmental impact to ethical considerations.

Supporters, particularly tech companies, argue that AI can personalize learning to meet each child's unique needs. The Genius Group, for instance, has developed AI avatars, like "Alan Turing," (Genius Group, 2024a) designed to adapt to individual learning paces, preferences, and challenges. Through platforms like "PopBots," (Williams et al., 2019) which uses robotics to introduce AI concepts to children as young as four, AI becomes not only a tool for customized learning but also an early step toward AI literacy and critical thinking. Advocates believe that early exposure to AI will equip the next generation with essential technological understanding, preparing them for a future increasingly shaped by AI.

Some education specialists view AI as a solution to the challenge of individualized education. In early childhood, finding the "sweet spot" of learning for each child is crucial but difficult in a classroom setting. AI-driven tools could help educators by providing activities tailored to each child's skill level, fostering engagement and critical thinking as tasks evolve with the child's progress (Su & AbstractNowadays, 2023).

Despite these promises, concerns persist, particularly among parents and child safety advocates. The Child Rescue Coalition highlights risks of AI exposing young users to inappropriate content. AI models, even with filters, can inadvertently display harmful material, as seen with conversational platforms like Character AI (Sina, 2024). Moreover, AI could facilitate troubling risks, such as deepfake impersonations and AI-enhanced grooming tactics (Sina, 2024). These issues underscore the need for rigorous oversight, especially as younger children lack the maturity to recognize manipulation, making regulation essential.

Different perspectives among participants further complicate AI's role in early childhood education. Educators, while cautiously optimistic about AI's potential, stress the importance of boundaries and training to ensure it complements traditional teaching without replacing

essential face-to-face interactions. Many teachers face challenges in incorporating AI due to limited guidelines and training (Child Rescue Coalition, 2024).

Parents, meanwhile, are divided, with some interested in AI's developmental support and others wary of its impact on social skills and mental health (Famly, n.d.). Concerns about data privacy and AI's ethical implications loom large, as parents question the effects of children forming bonds with AI characters rather than human connections.

Tech companies, strong proponents of AI in education, emphasize AI's ability to modernize learning by personalizing experiences. Companies like Genius Group argue that AI can support early learning while fostering AI literacy (Genius Group, 2024b). However, skepticism arises over whether profit motives overshadow child welfare, as rapid technological advances often outpace regulatory frameworks, leading to oversight gaps.

This sociotechnical research explores these perspectives, aiming to understand how each participant's priorities shape their stance on AI in early education. Educators focus on pedagogical impact, parents prioritize safety and ethics, and tech companies emphasize innovation and personalization. The sociotechnical research reveals how diverse values and concerns drive differing stances on AI in early childhood education, highlighting the need for balanced, informed regulation that aligns technological innovation with developmental and ethical priorities crucial for young learners.

References

AI in early childhood education: Famly. RSS. (n.d.). https://www.famly.co/us/blog/ai-in-early-childhood-education

Chakraborty, T., & Masud, S. (2023, June 7). Judging the creative prowess of AI. *Nature News*. https://www.nature.com/articles/s42256-023-00664-y

Cseh, G. M., & Jeffries, K. K. (2019). A scattered CAT: A critical evaluation of the consensual assessment technique for creativity research. *American Psychological Association*. https://psycnet.apa.org/record/2019-20312-005

Fan, L., Wang, G., Jiang, Y., Mandlekar, A., Yang, Y., Zhu, H., Tang, A., Huang, D.-A., Zhu, Y., & Anandkumar, A. (2022, November 22). Minedojo: Building open-ended embodied agents with internet-scale knowledge. *arXiv.org*. https://arxiv.org/abs/2206.08853

Genius Group appoints Alan Turing AI as chief AI officer. (2024, April 24). *Genius Group Limited*. https://ir.geniusgroup.net/news-events/press-releases/detail/131/genius-group-appoints-alan-turing-ai-as-chief-ai-officer

Genius Group launches AI Avatar tutor team on GeniusU. (2024, February 14). *Genius Group Limited*. https://ir.geniusgroup.net/news-events/press-releases/detail/119/genius-group-launches-ai-avatar-tutor-team-on-geniusu

Kolbytn, M. (n.d.). mindcraft. *GitHub*. https://github.com/kolbytn/mindcraft

Nottingham, K., Ammanabrolu, P., Suhr, A., Choi, Y., Hajishirzi, H., Singh, S., & Fox, R. (2023, April 27). DO embodied agents dream of pixelated sheep: Embodied decision making using language guided World modelling. *arXiv.org*. https://arxiv.org/abs/2301.12050

Sina. (2024, July 18). Is character AI safe or a safety risk? A guide for parents. *Screentime Monitoring | Safes Parental Control App*. https://www.safes.so/blogs/is-character-ai-safe/

Su, J., & AbstractNowadays. (2023, January 11). Artificial Intelligence (AI) literacy in early childhood education: The challenges and opportunities. *Computers and Education: Artificial Intelligence*. https://www.sciencedirect.com/science/article/pii/S2666920X23000036

The Dark Side of AI: Risks to children. (2024, June 18). *Child Rescue Coalition*. https://childrescuecoalition.org/educations/the-dark-side-of-ai-risks-to-children/

Torrance, E. P. (n.d.). Torrance Tests of Creative Thinking. *American Psychological Association*. https://psycnet.apa.org/record/9999-05532-000?doi=1

Williams, R., Park, H. W., Oh, L., & Breazeal, C. (2019). Popbots: Designing an artificial

intelligence curriculum for early childhood education. *Proceedings of the AAAI Conference on Artificial Intelligence*.

https://ojs.aaai.org/index.php/AAAI/article/view/5040