Learning How Large Language Models Work Through Games

ALLISON CHEN, Princeton University, USA

ISABELLA PU, Massachusetts Institute of Technology, USA

While artificial intelligence (AI) technology is becoming increasingly popular, its underlying mechanisms tend to remain opaque to most people. To address this gap, the field of AI literacy aims to develop various resources to teach people how AI systems function. Here we contribute to this line of work by proposing two games demonstrating some basic mechanisms behind large language models (LLMs), a popular type of AI system. The first game, *Learn Like an LLM*, aims to convey that LLMs are trained to predict sequences of text based on a particular dataset. The second game, *Tag-Team Text Generation*, focuses on teaching that LLMs generate text one word at a time, using both probability and randomness. While these ideas proposed are still in early stages and would benefit greatly from further discussion, we hope they can contribute to using game-based learning to teach about complex AI systems like LLMs.

Additional Key Words and Phrases: Artificial Intelligence, AI Literacy, Large Language Models, Games

ACM Reference Format:

Allison Chen and Isabella Pu. 2025. Learning How Large Language Models Work Through Games. In *Proceedings of CHI Conference on Human Factors in Computing Systems (CHI '26).* ACM, New York, NY, USA, 9 pages. https://doi.org/X

1 Introduction

As artificial intelligence (AI) technology becomes increasingly prevalent, there is a need for people to develop an accurate understanding about these systems, including how they work and how to use them. Since ChatGPT released in 2022, AI has affected many spheres of life, such as the economy [18, 41], politics and policy [3], and daily life [5, 10, 28, 39]. Despite its growing presence, many people have limited access to the resources and training needed to develop an accurate understanding about the AI technology. This gap is further widened by the common use of anthropomorphic, or human-like, language used directly by AI systems or used to describe AI systems [15, 35]. While this use of anthropomorphic language may make complex AI systems *feel* more approachable, it simultaneously masks the technical and algorithmic nature of these systems. As a result, this can lead to increased risk of people developing emotional attachments and overestimating the system's capabilities [1, 11, 21].

To help people develop a better understanding about AI technology, the field of AI literacy seeks to develop educational materials and resources that combat misleading anthropomorphic narratives about AI [23, 25, 30]. Building upon various forms of literacy (e.g., digital, computational, scientific, and data literacy), AI literacy encompasses many aspects of how people interact with and relate to AI systems, such as what it can do, how it works, how it should be used, and how people perceive these systems [25]. Understanding how these systems work and what they can and cannot do is important for making informed decisions, as AI now plays a role in areas ranging from financial investments to regulation and policy to everyday use at work or in the home [17]. Familiarity with the technical components of AI

Authors' Contact Information: Allison Chen, allisonchen@princeton.edu, Princeton University, Princeton, New Jersey, USA; Isabella Pu, ipu@media.mit.edu, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

systems can help reduce the sense of "mystery" surrounding AI technology and reduce people's overreliance on these systems [1, 40, 42]. Though much of the work on AI literacy for conveying technical concepts has focused primarily on more formal academic classroom settings [8, 20, 37], a growing number of works are recognizing the value of developing AI literacy materials for informal settings, such as museums or other public spaces [13, 14, 24, 33]. These works recognize the benefit of informal learning, as such settings can be more engaging, allow parents and children to learn together, and make the information more accessible.

In this work, we build upon existing works promoting AI literacy through informal learning settings by outlining novel ideas for two games. These games aim to intuitively convey fundamental ideas behind how modern AI systems, specifically large language models (LLMs), work. We chose to focus on LLMs because they are a popular AI system and many people use them, either directly through chat interfaces or indirectly through other software products (e.g., Google search) [4, 22]. Prior AI literacy works tended to focus on foundational concepts like machine learning or neural networks [7, 36], but it is becoming increasingly important to also develop materials about specific AI systems that people may encounter frequently, such as LLMs [2, 27]. While there are a few works that focus on generative AI systems like LLMs, these primarily aim to explain the behaviors or impact of the systems, not necessarily how they work [6, 20, 27]. On the other hand, our proposal is centered around two games designed to convey technical concepts about the mechanisms of LLMs produce text.

We decided to use games because they have been shown to be effective learning tools [32, 38] for adults as well as children [9]. First, games can be motivating, engaging, and adaptive to different learning styles [32]. They further provide low-stakes opportunities for experimentation and feedback, creating a safe space for failure that is important to effective learning [12, 43]. Games also help foster autonomy and curiosity, making abstract or intimidating concepts like AI more approachable through playful exploration and immediate feedback [19, 29, 34]. Specifically in AI literacy, games have been increasingly employed and shown to effectively teach AI-related concepts [14, 16, 31, 33] while additionally providing positive affective experiences [26, 31]. Here, we build upon game-based learning to develop AI literacy materials regarding specifically how LLMs work.

We propose two games with the goal of filling the gap in the public's familiarity with LLM systems and understanding of how they actually work. The first game focuses on how LLMs learn to generate text via predicting sequences of words from a dataset, and the second focuses on the how LLMs generate text one word at a time and by using both probabilities and randomness. In first game, *Learn Like an LLM*, players are tasked with guessing sequences of shapes that belong to a hidden set, similar to how LLMs are tasked to predict sequences of text from the training dataset. The primary learning objective is for players to understand that the LLM training process is more analogous to pattern recognition than it is to how people learn to use language to communicate ideas. In the second game, *Tag-Team Text Generation*, players take turns with a "computer player" to generate text responses, one word at a time. Here, the primary learning objective is for players to recognize that LLMs do not "think about ideas" then find the words to convey them like people do. Rather, LLMs generate text by predicting one word[1] at a time, based on both probability predictions and randomness. We note that these games do not aim to provide people with a *complete* technical understanding of LLMs. Rather, they aim to give players an *intuition* about specific aspects of LLMs that lead them to say *"This isn't how I thought these technologies worked!"* or *"This is not so complicated, I can understand this!"*

[1]Truly, it is a sub-word token, but for the purposes of our game, we do not go into this detail.

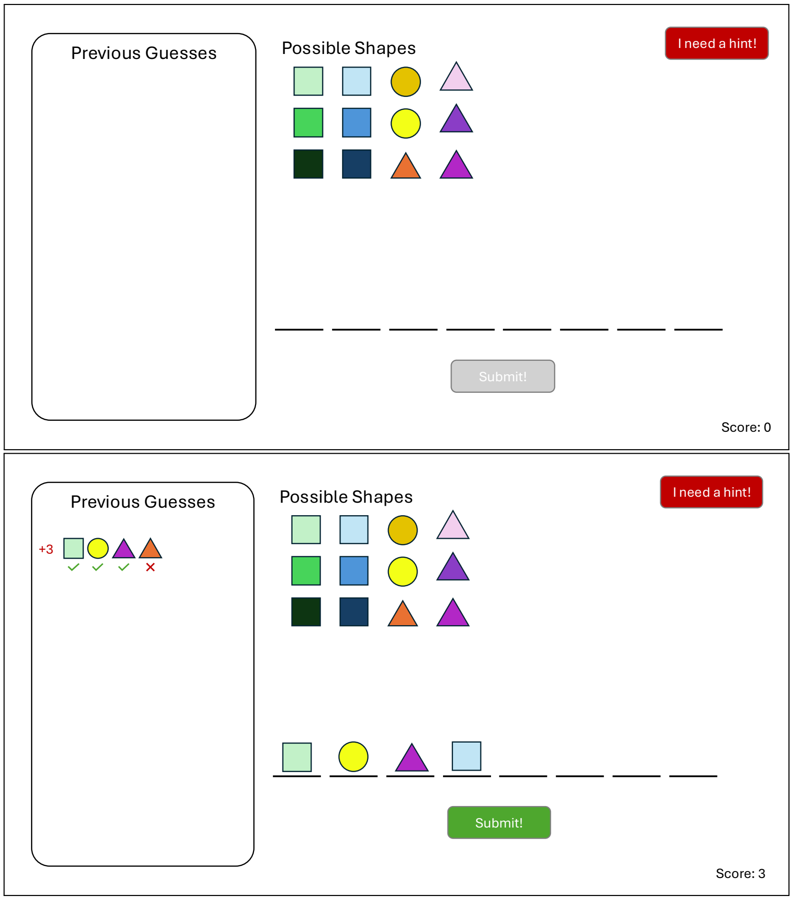

Fig. 1. Two example interfaces of Game 1 (*Learn Like an LLM*). **Top**: This is the interface when the player starts. The set of possible shapes is always at the top of the screen. Once the player selects at least 4 shapes, the Submit button becomes clickable. **Bottom**: A possible interface after the player submits one sequence and is constructing the second. The submitted sequene is on the left side of the screen. The points earned from the first sequence is to the left of the sequence, the color of the points denotes whether the shape was in the hidden set, and the validity of each shape is denoted by the checkmarks and X's under each shape.

2 Game 1: *Learn Like an LLM*

The main goal of *Learn Like an LLM* is to help people understand that LLMs are optimized to predict words that are likely to come next given a sequence of text and it is heavily dependent on a particular dataset. To convey this idea, players are given a “vocabulary” of 10-12 simple shapes of different colors. Then they are given the following instructions:

(1) There is a hidden set of sequences of shapes. Each sequence is between 4-8 shapes long, and they follow some patterns or rules. Your goal is to guess as many sequences in the hidden set as possible!
(2) You can build a sequence by selecting one shape at a time, and you may submit your guess for a sequence once you have selected at least 4 shapes. Or you may continue to build the sequence up to 8 shapes long.
(3) When the sequence is submitted, you may receive points for each shape in the sequence and for the whole sequence. You will receive one point for each shape that is valid given the preceding shape (or if it is the first

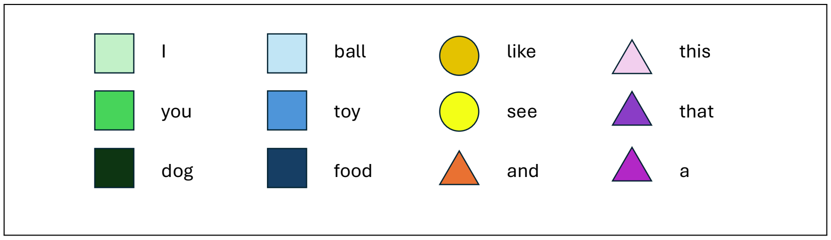

Fig. 2. Possible mapping of shapes to words for *Learn Like an LLM*. Based on this mapping, the previously submitted sequence in Fig. 1 (bottom) maps to "I see a and" and the sequence in progress maps to "I see a ball".

shape, whether it is a valid first shape). Second, if all shapes are valid *and* the sequence was in the hidden set, you will receive additional 3 points for guessing the sequence. Note, not all valid sequences are in the hidden set, *but* this can still be helpful information to help you deduce the rules governing the sequences.

(4) If you get stuck, you may click on the "I need a hint!" button, and this will reveal a sequence from the hidden set. However, this counts as one of your guesses and you will not receive any points from it.

(5) You have 10 tries to guess as many sequences as you can!

After reading the instructions, the player will see the game interface shown in Fig. 1 (top) where they can guess which sequences of shapes belong to the hidden set. The current sequence will be built on the dashed lines and the "Submit!" button is valid once there is a minimum of four shapes selected. When the player submits a sequence, it will appear on the left hand side with the points received, as shown in Fig. 1 (bottom). The "+3" denotes the player earned 3 points from that sequence, one for each shape with a green checkmark under it indicating that it was valid. The red color of the "+3" indicates that the sequence was not in the hidden set. The game ends after the player submits 10 guesses. Unbeknownst to the player, each shape actually represents an English word and the sequences are based on grammatically correct sentences[2]. An example shape-to-word mapping can be found in Fig. 2. Because the sequences mirror English grammar, the goal is that players will pick up on patterns, such as how certain shapes tend to follow others.

After the game ends, this shape-to-word mapping will be revealed and the player will receive a debriefing on how the game mirrors the way LLMs are trained. Primarily, the hidden set reflects LLM training data where LLMs are trained by predicting the next word given a sequence of text. Similarly, in this game, the players have a sequence of shapes that they have to predict one shape at a time. Further, at the start of the game, the shapes had no semantic meaning or relation to one another for the player. However, over the course of receiving feedback for correct and incorrect guesses, the player may learn associations between certain shapes without necessarily ascribing meaning to the shapes. This mirrors how before training, words do not carry semantic meanings to LLMs, but the LLM can model statistical patterns between specific words. We hope this point raises the question for the player to ponder and discuss: even if LLMs can robustly model the statistical relations of words, does this mean it knows the words' meaning? Lastly, the hidden set of sequences did not contain *all* grammatically correct combinations of the shapes, which meant that the player is penalized for valid English sequences beyond those in the set. This is similar to how LLMs have non-exhaustive training datasets and can be constrained by their training data. Practically, this can lead LLMs to pick up biases in the training data. In summary, *Learn Like an LLM* aims to players how LLMs are trained to generate text by predicting sequences of text in a particular dataset and are penalized for wrong predictions and rewarded for correct predictions.

[2] Barring conjugations for tense and plurality.

3 Game 2: *Tag-Team Text Generation*

The second game, *Tag-Team Text Generation*, is a creative and non-competitive activity that aims to convey how LLMs generate text one word at a time[3] and how prediction utilizes both probabilities and randomness. In this game, the player will alternate with the computer to simulate how LLMs generate responses: one word at a time and through a two-step process using probabilities and randomness.

The player starts by selecting a fun prompt to answer, such as one of the following:

- What would Cinderella's godmother give her if she lived in 2026?
- Describe a sunset on the beach from the perspective of a baby turtle who just hatched.
- Create your own fantasy creature, describing what it looks like and what it likes to do.
- Where is the best place to live—real or imaginary—and why?
- If you were a ruler of a country, what would be your national food and why?

Alternatively, players can also submit their own prompt. Then to build the response, the player and computer will alternate between generating a set of probable words and selecting a word from the generated set.

To start, the computer will generate and display a set of 5 possible words with their corresponding probabilities, such as shown in Fig. 3 (top). The probabilities do not necessarily need to add up to 100% as these words only represent a *subset* of all possible options. The player then selects any word they like—it could be random, based on what they think is funny, etc. They do not need to choose the one that is most probable, and in fact, varying their selections will likely be more entertaining. Once the player selects the word, it is added to the response as shown in Fig. 3 for the word "Magnificent!"

To generate the next word, the roles are flipped. Now, the player will provide the computer with a small set of possible words that the computer will randomly select from. To do so, the player can either choose to come up with three words they think are most probable *or* select three from a larger pool of possible words if they feel this is too cumbersome. If they select the latter option, the player is presented with 10 possible words (without any probabilities), as shown in Fig. 3 (bottom). The computer will then take the three submitted words and randomly select one of them as the next word in the response.

These two steps—where the player selects a final word or creates a pool of probable words—repeats until the player is satisfied with their response and can submit it. Upon submitting, the player can see other people's responses to the same prompt; this will be especially conducive for an online setting as well as in museums and public spaces. As an additional reward, we can add the option for players to submit their prompt and generated response to a text-to-image generation model and print out the resulting image.

After the player is done, they will receive a debriefing of how this game reflects the text generation mechanisms of LLMs. First, players experienced needing to generate only one word at a time; this limits how much they could "plan ahead" on what to write because they were limited by the computer's estimates of which words were most probable or by the computer's random selection. The players may have experienced hoping to take the response in one direction but the computer's selection was incompatible with their plan. Similarly, LLMs do not "think ahead" about what text to generate. Second, players experienced a two-step process: selecting a pool of possible words based on probabilities and choosing a word from that pool. LLMs undergo a similar process when they generate words: they first output how probable each word is, then employ an algorithm that employs randomness to select the final prediction for each word. This randomness is what allows for diverse responses to the same prompt. Conveying that randomness is a part of the

[3]As mentioned in the introduction, in practice, LLMs generate text one *token* at a time where tokens can be sub-words or full words, but we omit this detail in the game for brevity.

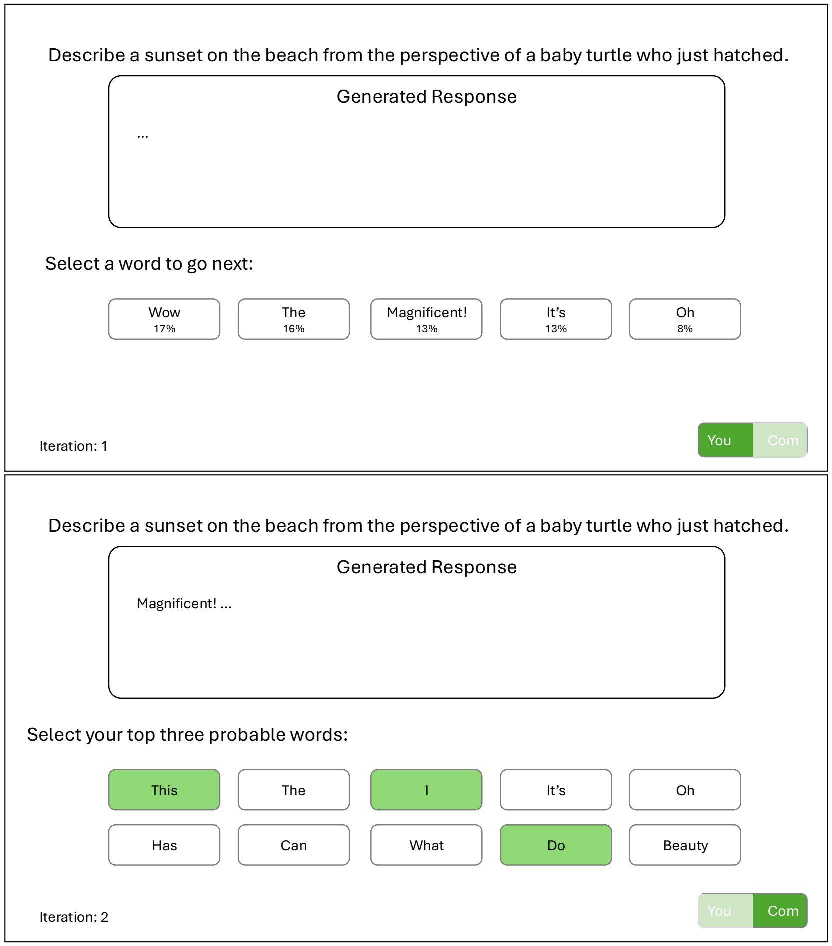

Fig. 3. Interface for *Tag-Team Text Generation*. **Top**: Interface when it is the player's turn to make the final word selection from a set of 5 words with their estimated probabilities. Here they select one. **Bottom**: Interface when it is the player's turn to generate a set of probable words. In this example, the player opted to select three from a pool of 10, rather than submitting their own. The computer will select one of the three selected words to add to the response.

text generation process further emphasizes the idea from *Learn Like an LLM* that LLMs do not generate text by coming up with ideas to then convey through text. In summary, *Tag-Team Text Generation* aims to demonstrate to players that LLMs predict text one word at a time through a process that involved both probabilities and randomness.

4 Discussion

4.1 Summary

In this work, we propose two games, *Learn Like an LLM* and *Tag-Team Text Generation* that aim to convey technical ideas about how LLMs work, specifically about the training procedure and text generation process respectively. In *Learn Like an LLM*, players are taught that LLMs are trained to produce text much like one can predict sequences of shapes through pattern recognition. The game is meant to reinforce that LLMs do not learn to generate text the same way people learn to use words to communicate ideas, but rather that LLMs operate based on pattern recognition and modeling. In the second game, *Tag-Team Text Generation*, players learn that the LLM text generation process can be broken down into two steps: using probabilities to select a subset of words that are likely to come next and selecting one of them, perhaps randomly but importantly, not necessarily by selecting the word with the highest probability. This game is meant to reveal that the generation procedure is not as sophisticated as companies may suggest, but is based on probabilities and randomness.

4.2 Limitations

While these games aim to shed light on some aspects of how LLMs work, we acknowledge a few limitations. First, the games still abstract away certain details and are not perfect reflections of how LLMs are trained and predict text. For example, LLMs are provided a sequence of text as an input and are tasked to generate one word, but in *Learn Like an LLM*, players are building up sequences from scratch. With *Tag-Team Text Generation*, it seems that LLMs output only a few words that are probable, but in reality, they output a probability distribution over the entire vocabulary. A second limitation is that there are many aspects of LLMs that these games do not touch on. For example, they do not shed light on how words are represented as vectors, how text gets split up into units called tokens, or how, in practice, many LLMs are likely to use more sophisticated sampling techniques to select the final word. We hope that future work can focus on developing games and activities to address these topics. Lastly, it is possible for people to play these games without ever connecting them to LLMs. While this may still provide a fun and engaging experience, we hope to improve the games by highlighting how they are connected to LLMs before the debriefing at the end.

4.3 Conclusion

In this work, we outline early stage ideas for two games conveying how LLMs work: one focused on the training procedure and the other on text generation. We hope that the CHI 2026 Workshop on Data Literacy will be a helpful opportunity to receive feedback and actively iterate on these ideas to address the aforementioned limitations, better convey intuitions about how LLMs work, and improve the gameplay experience overall. We additionally hope that this work will open up conversations for novel and engaging ways of using games to teach people—children and adults alike—about how LLMs and other modern AI systems work through engaging and fun activities!

References

[1] Rena Alasgarova and Jeyhun Rzayev. 2025. Managing Anthropomorphism in Student–AI Interaction. *International Journal of Information Technology and Education* 5, 1 (2025), 139–160.

[2] Ravinithesh Annapureddy, Alessandro Fornaroli, and Daniel Gatica-Perez. 2025. Generative AI literacy: Twelve defining competencies. *Digital Government: Research and Practice* 6, 1 (2025), 1–21.

[3] Eric Best, Pedro Robles, and Daniel J Mallinson. 2024. The Future of AI Politics, Policy, and Business. *Business and Politics* 26, 2 (2024), 171–179.

[4] Michelle Brachman, Amina El-Ashry, Casey Dugan, and Werner Geyer. 2025. Current and future use of large language models for knowledge work. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–24.

[5] Petter Bae Brandtzaeg, Asbjørn Følstad, and Marita Skjuve. 2025. Emerging AI Individualism: How Young People Integrate Social AI into Everyday Life. *Communication and Change* 1, 1 (2025), 11.

[6] Huajie Jay Cao, Hee Rin Lee, and Wei Peng. 2025. Empowering Adults with AI Literacy: Using Short Videos to Transform Understanding and Harness Fear for Critical Thinking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–8.

[7] Michelle Carney, Barron Webster, Irene Alvarado, Kyle Phillips, Noura Howell, Jordan Griffith, Jonas Jongejan, Amit Pitaru, and Alexander Chen. 2020. Teachable machine: Approachable Web-based tool for exploring machine learning classification. In *Extended abstracts of the 2020 CHI conference on human factors in computing systems*. 1–8.

[8] Lorena Casal-Otero, Alejandro Catala, Carmen Fernández-Morante, Maria Taboada, Beatriz Cebreiro, and Senén Barro. 2023. AI literacy in K-12: a systematic literature review. *International Journal of STEM Education* 10, 1 (2023), 29.

[9] Nathalie Charlier, Michela Ott, Bernd Remmele, and Nicola Whitton. 2012. Not just for children: game-based learning for older adults. In *6th European Conference on Games Based Learning, Cork, Ireland*. 102–108.

[10] Aaron Chatterji, Thomas Cunningham, David J Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. *How People Use ChatGPT*. Technical Report. National Bureau of Economic Research.

[11] Myra Cheng, Alicia DeVrio, Lisa Egede, Su Lin Blodgett, and Alexandra Olteanu. 2024. " I Am the One and Only, Your Cyber BFF": Understanding the Impact of GenAI Requires Understanding the Impact of Anthropomorphic AI. *arXiv preprint arXiv:2410.08526* (2024).

[12] Alan F. Chow, Kelly C. Woodford, and Jeanne Maes. 2011. Deal or No Deal: using games to improve student learning, retention and decision-making. *International Journal of Mathematical Education in Science and Technology* 42, 2 (March 2011), 259–264. doi:10.1080/0020739X.2010.519796

[13] Hasti Darabipourshiraz, Dev Ambani, and Duri Long. 2024. DataBites: An embodied and co-creative museum exhibit to foster children's understanding of supervised machine learning. In *Proceedings of the 16th Conference on Creativity & Cognition*. 550–555.

[14] Hasti Darabipourshiraz, Lily Murakami Ng, Grace Wang, Sophie Rollins, and Duri Long. 2026. AI Unplugged: Exploring Pathways from Physical Simulation to Conceptualization of AI Reasoning Processes. *ACM Transactions on Computing Education* (2026).

[15] Alicia DeVrio, Myra Cheng, Lisa Egede, Alexandra Olteanu, and Su Lin Blodgett. 2025. A Taxonomy of Linguistic Expressions That Contribute To Anthropomorphism of Language Technologies. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.

[16] Xiaoxue Du, Zhichun Liu, Xi Wang, and Qiping Tang. 2024. Fostering AI Literacy through Interactive Game-Based Learning: A Case Study on Enhancing Algorithmic Thinking Skills. In *Society for Information Technology & Teacher Education International Conference*. Association for the Advancement of Computing in Education (AACE), 333–338.

[17] Fariba Fard and Suliman Hawamdeh. 2026. The Growing Need for AI Literacy in Business and Daily Life. In *AI-Driven Revolution: Transforming the Business Landscape*. World Scientific, 21–36.

[18] Mirko Farina, Xiao Yu, and Andrea Lavazza. 2025. Ethical Considerations and Policy Interventions Concerning the Impact of Generative AI Tools in the Economy and in Society. *AI and Ethics* 5, 1 (2025), 737–745.

[19] Michail Giannakos, Iro Voulgari, Sofia Papavlasopoulou, Zacharoula Papamitsiou, and Georgios Yannakakis. 2020. Games for Artificial Intelligence and Machine Learning Education: Review and Perspectives. In *Non-Formal and Informal Science Learning in the ICT Era*, Michail Giannakos (Ed.). Springer Singapore, Singapore, 117–133. doi:10.1007/978-981-15-6747-6_7 Series Title: Lecture Notes in Educational Technology.

[20] Matthias Carl Laupichler, Alexandra Aster, Jana Schirch, and Tobias Raupach. 2022. Artificial intelligence literacy in higher and adult education: A scoping literature review. *Computers and Education: Artificial Intelligence* 3 (2022), 100101.

[21] Jiaxuan Li, Mingxing Han, and Qinjian Yuan. 2025. The Surprising Paradox of AI Literacy: How Lower AI Literacy Can Lead to Higher Acceptance. *International Journal of Human–Computer Interaction* (2025), 1–21.

[22] Weixin Liang, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, Sheng Liu, Siyu He, Zhi Huang, et al. 2025. Quantifying large language model usage in scientific papers. *Nature Human Behaviour* (2025), 1–11.

[23] Tomáš Lintner. 2024. A systematic review of AI literacy scales. *npj Science of Learning* 9, 1 (2024), 50.

[24] Duri Long, Takeria Blunt, and Brian Magerko. 2021. Co-designing AI literacy exhibits for informal learning spaces. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–35.

[25] Duri Long and Brian Magerko. 2020. What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–16.

[26] Duri Long, Anthony Teachey, and Brian Magerko. 2022. Family learning talk in AI literacy learning activities. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–20.

[27] Qianou Ma, Anika Jain, Jini Kim, Megan Chai, and Geoff Kaufman. 2025. ImaginAItion: Promoting Generative AI Literacy Through Game-Based Learning. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–9.

[28] Brandon T McDaniel, Fayika F Nova, and Jessica A Pater. 2025. Artificial Intelligence in Everyday Family Life: Issues, Applications, and Implications. *Family Relations* 74, 3 (2025), 1049–1055.

[29] Kalpana Nand, Nilufar Baghaei, John Casey, Bashar Barmada, Farhad Mehdipour, and Hai-Ning Liang. 2019. Engaging children with educational content via Gamification. *Smart Learning Environments* 6, 1 (July 2019), 6. doi:10.1186/s40561-019-0085-2

[30] Davy Tsz Kit Ng, Jac Ka Lok Leung, Samuel Kai Wah Chu, and Maggie Shen Qiao. 2021. Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence* 2 (2021), 100041.

[31] Davy Tsz Kit Ng, Chen Xinyu, Jac Ka Lok Leung, and Samuel Kai Wah Chu. 2024. Fostering students' AI literacy development through educational games: AI knowledge, affective and cognitive engagement. *Journal of computer assisted learning* 40, 5 (2024), 2049–2064.

[32] Jan L Plass, Bruce D Homer, and Charles K Kinzer. 2015. Foundations of game-based learning. *Educational psychologist* 50, 4 (2015), 258–283.

[33] Sophie Rollins, Katherine Hancock, Jasmin Ali-Diaz, Nyssa Shahdadpuri, and Duri Long. 2024. Knowledge Net: Fostering Children's Understanding of Knowledge Representations Through Creative Making and Embodied Interaction in a Museum Exhibit. In *Proceedings of the 16th Conference on Creativity & Cognition*. 470–475.

[34] Silmara Rondon, Fernanda Chiarion Sassi, and Claudia Regina Furquim De Andrade. 2013. Computer game-based and traditional learning method: a comparison regarding students' knowledge retention. *BMC Medical Education* 13, 1 (Dec. 2013), 30. doi:10.1186/1472-6920-13-30

[35] Arleen Salles, Kathinka Evers, and Michele Farisco. 2020. Anthropomorphism in AI. *AJOB neuroscience* 11, 2 (2020), 88–95.

[36] Daniel Smilkov, Shan Carter, D Sculley, Fernanda B Viégas, and Martin Wattenberg. 2017. Direct-manipulation visualization of deep networks. *arXiv preprint arXiv:1708.03788* (2017).

[37] Jiahong Su, Davy Tsz Kit Ng, and Samuel Kai Wah Chu. 2023. Artificial intelligence (AI) literacy in early childhood education: The challenges and opportunities. *Computers and Education: Artificial Intelligence* 4 (2023), 100124.

[38] Sigmund Tobias, J Dexter Fletcher, and Alexander P Wind. 2013. Game-based learning. *Handbook of research on educational communications and technology* (2013), 485–503.

[39] Herminio Pabón Trujillo, Juan Carlos Castillo, Angie Dayana Rangel, Oscar Fabián Patiño, Ronald Edinxon Angarita Bautista, and Jenny Zulay Cáceres Cárdenas. 2025. The Global Impact of AI on Workplace Safety, Opportunities and Challenges for the Future of Work. *Ciencia Latina Revista Científica Multidisciplinar* 9, 2 (2025), 7500–7513.

[40] Stephanie M Tully, Chiara Longoni, and Gil Appel. 2025. Lower artificial intelligence literacy predicts greater AI receptivity. *Journal of Marketing* (2025), 00222429251314491.

[41] Qiang Wang, Tingting Sun, and Rongrong Li. 2025. Does Artificial Intelligence (AI) Enhance Green Economy Efficiency? The Role of Green Finance, Trade Openness, and R&D Investment. *Humanities and Social Sciences Communications* 12, 1 (2025), 1–22.

[42] Geert Wood, Elena Nuñez Castellar, and Wijnand IJsselsteijn. 2025. An Exploratory Study Into the Impact of AI Literacy Training on Anthropomorphism and Trust in Conversational AI. In *International Conference on Human-Computer Interaction*. Springer, 301–322.

[43] Jialing Zeng, Sophie Parks, and Junjie Shang. 2020. To learn scientifically, effectively, and enjoyably: A review of educational games. *Human Behavior and Emerging Technologies* 2, 2 (2020), 186–195. doi:10.1002/hbe2.188

Received 10 February 2026