Games of Representation: Developing Card-Based Activities to Teach About Representation and Bias in AI Datasets

Katherine S. Moore[1], **Helen Zhang**[2], **Irene Lee**[3]

[1] MIT STEP Lab, 600 Technology Square NE49-3021, Cambridge MA 02139, USA

[2] Lynch School of Education and Human Development, Boston College, Chestnut Hill MA 02459, USA

[3] New Mexico State University, Las Cruces, NM 88003, USA

ksmoore@mit.edu[1], zhangzm@bc.edu[2], ialee@nmsu.edu[3]

Abstract

Adolescents struggle to understand bias and representation in AI, particularly the concept of how datasets used in machine learning can be representative of populations or not. Within early experiments in teaching about investigating bias in AI systems using the Developing AI Literacy (DAILy) curriculum, we observed participating youth struggling to understand what it meant to be represented in the output of AI tools. For example, when using Google Image Search with prompts such as "physicist" and "outdoor recreation," participating youth did not understand the question, "Are you represented in this outcome?" We saw an opportunity to address this challenge using the Kapor Foundation's Responsible AI and Tech Justice Guide. Drawing insights from three of the six core components of the framework presented in the guide, we developed Games of Representation (GR), a series of three card-based activities using SET game cards to teach concepts of population, sample, dataset, and representation. Through game play, players manipulate datasets, role-play as stakeholders with competing interests, and explore real-world scenarios where representation matters. The GR games and corresponding guide provide flexible, hands-on activities for guided playful conversations with adolescents about AI ethics. This work contributes practical resources for K-12 AI education, addressing challenges in youth understanding of statistical bias and stakeholder influence in AI development.

Extended version — Everyday-AI.org

Introduction

The rapid integration of artificial intelligence (AI) into daily life has created an urgent need for AI literacy education at the K-12 level. Adolescents encounter AI-powered technologies through search engines, social media algorithms, and emerging applications in education, healthcare and everyday technologies. As these systems increasingly influence access to opportunities and resources, developing understanding of how AI works—and how it can amplify societal biases—has become essential for informed citizenship and ethical participation in a technology-mediated world.

A common challenge in AI literacy education involves cementing the understanding of the relationship between data representation and AI predictions. Our observations from implementations of the Developing AI Literacy (DAILy) curriculum (Lee et. al. 2021; Zhang et. al. 2023; Zhang et. al. 2024) revealed a particular manifestation of this challenge. For example, frequently, when participating students examined Google Image Search results for terms like "physicist," many were often able to indicate that the results seemed biased. However, when asked "Are you represented in this dataset?" students consistently struggled to understand what representation meant in the context of datasets and AI systems. Many could not articulate what features or characteristics might be relevant for determining representation, nor could many connect the concept of being represented in training data to the biased outcomes they observed from the image search.

This lack of understanding representation may lead to consequent challenges when learners try to make sense of a) how datasets shape AI behavior, b) what it means for datasets to be representative of larger populations, and c) how human decisions about data collection and curation influence AI predictions.

To address these challenges, we developed Games of Representation (GR), a series of three card-based activities designed to make abstract concepts of dataset representation tangible and accessible through hands-on gameplay. Leveraging cards from the SET game[1], an easily accessible game, the games engage learners in actively constructing, manipulating, and analyzing datasets while role-playing as stakeholders with competing interests in AI development.

[1] The SET game deck was created by Marsha Jean Falco for Set Enterprises © 2019 Cannei LLC. All rights reserved.

This paper describes the design of these games, with a focus on how hands-on, low-tech approaches can address critical gaps in AI literacy education. We 1) present the background and design goals underlying the games, 2) provide detailed descriptions of the three-game sequence, and 3) discuss key design features that distinguish this approach from existing AI education resources.

GR aims to offer educators, families and caregivers, interested in talking with their students and adolescents about foundational concepts in AI, a promising pedagogical tool and guide for developing adolescents' understanding of representation, bias, and stakeholder influence in AI while providing accessible entry points for rich discussions about AI ethics across diverse contexts including informal education, intergenerational and care-giver settings.

Design Background: A Need Emerged from Prior Work

GR emerged from implementation challenges identified by teachers, field investigators, and members of our research team during the deployment of the DAILy curriculum across diverse educational contexts, including formal classroom and informal educational settings (e.g., AI summer camps). The DAILy curriculum aims to develop critical AI literacy among middle school age learners through hands-on activities that examine how AI systems impact daily life. One key lesson in the curriculum, *Investigating Bias*, engages learners in interrogating AI systems by running simple queries in tools like Google Image Search and critically assessing the results returned.

In a typical implementation of the activity for this lesson, learners search for terms such as "physicist" or "outdoor recreation" and examine the demographic representation in the image results. The activity guides learners through structured reflection questions: What did you investigate? What did you find out? Who is represented? Who is left out? Was the result accurate? Is it fair? Why do you think this happened? What are the implications? What could be done about it?

Through implementation with over 2,000 students across six U.S. states between the years 2020-2024 a semi-consistent pattern emerged: while participating adolescents could frequently answer the first two questions about their investigation and findings, they struggled with questions about representation. They had difficulty understanding what it meant for people to be "represented" or "left out" of search results and could not articulate the relationship between biased outcomes and underlying dataset composition. Even when the question was reframed as "do you see yourself or people like you in the search results?" participants could respond about their personal experience, but did not grasp broader concepts of representation in the context of datasets and AI training processes.

This revealed a fundamental challenge; adolescents may lack understanding of how datasets shape AI behavior and what it means for datasets to be representative or non-representative of larger populations. In the light from this observed gap, our team saw an opportunity in an existing AI literacy framework, which guided our design goals.

Game Design Goals

The design goals for GR were informed by three core components from the *Kapor Foundation's Responsible AI and Tech Justice Guide*, which provides a framework for developing comprehensive AI literacy that addresses both technical and ethical dimensions of AI systems. The Kapor Foundation's approach aligns closely with our understanding of the challenges identified in the DAILy curriculum implementation. Where traditional AI education often focuses primarily on technical skills or algorithmic understanding, the Tech Justice Guide emphasizes the social contexts and power dynamics that shape AI development—precisely the conceptual gaps we observed in participating adolescents' responses to questions about representation and bias. The framework's integration of ecosystem understanding, critical thinking, and harm mitigation offered a comprehensive approach to addressing learners' need for frameworks that connect technical AI concepts to their social and ethical implications. We developed GR to address the following select core components (#1, #5, and #6) of the Kapor Foundation's Tech Justice Guide. Components #2-#4 were out of reach for the scope of this work as they touch on data security and privacy, the exploration of societal impacts, and the interrogation of one's personal use of technologies; all of which were beyond the scope and specific aims of this project.

Core Component #1: Examine the AI technology creation ecosystem

Our goal was to help adolescents understand how stakeholders and their competing interests impact dataset creation processes. We wanted to foster the skill of examining whose interests are served by particular dataset compositions and how power dynamics influence which voices and perspectives are included or excluded in the creation and curation of AI training data. This goal extends beyond technical understanding to encompass critical analysis of the social contexts in which AI are developed and deployed.

Core Component #5: Build a critical lens in the collection, usage, analysis, interpretation, and reporting of data

We aimed to engage learners in a hands-on practice of inspecting datasets and analyzing their contents. Ultimately, our goal was for players to develop the skill of describing

which features of a population are represented and not represented in datasets, while understanding how stakeholder roles and goals influence data manipulation decisions. Thus, through role-playing as stakeholders with specific objectives, players experience how competing interests shape dataset composition and learn to critically evaluate the representativeness of samples relative to larger populations.

Core Component #6: Learn how to minimize, eliminate, and mitigate harm and injustice caused by AI technologies

We sought to develop learners' capacity to analyze datasets for their degree of representativeness and to understand how these characteristics affect AI model predictions. A key goal involved helping learners recognize the contextual nature of representation—when being included in datasets might be beneficial versus harmful, and how participation in datasets can contribute to societal good or perpetuate injustice.

In conjunction with these core learning objectives were pedagogical goals to center accessibility and flexibility. We sought to create educational materials that could function effectively across diverse contexts—formal classrooms, informal afterschool programs, and intergenerational conversations. The games needed to be intellectually engaging for learners from 4th grade through high school while remaining accessible in resource-constrained settings without requiring technology infrastructure.

The Games of Representation

GR uses playing cards, as opposed to pieces on a board game or a puzzle, because doing so allows populations to be easily reconfigured. Also, the interpretation of the cards is straightforward due to their similarity to traditional card decks, which offers the familiarity of turn-taking and readily apparent card features (e.g., color, face value, suit).

GR consists of three sequential card-based activities using SET game cards, each building upon the previous to develop an increasingly sophisticated understanding of dataset representation and AI ethics. The games progress from concrete data exploration to strategic manipulation to real-world application, creating a scaffolded learning experience that makes abstract AI concepts more accessible.

Game 1: Investigating Samples

Game 1 introduces players to fundamental concepts of population, sample, and representation through guided exploration of a 16-card subset drawn from the complete SET deck (see Figure 1). Players work collaboratively to collaboratively to group cards by different features (color, shape, number, fill) and count how many cards represent each feature value (e.g., for the feature *color* the values are *red*, *green*, and *purple*) within the sample. This concrete

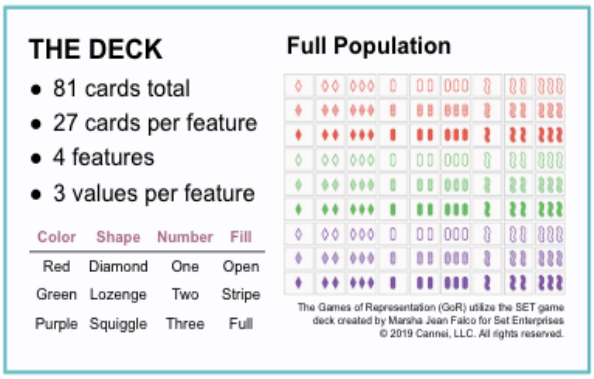

Figure 1: Illustrative summary of the SET Card Deck including the features and each features values (left) and a visualization of the full population (right).

counting and sorting helps players understand what constitutes a sample, identify which features are most and least represented, and begin to grasp the relationship between their sample and the larger population of all 81 cards.

This initial game emphasizes discovery and pattern recognition rather than competition, encouraging players to take turns forming different groupings and making observations about representation. Through systematic exploration, players develop a shared vocabulary around datasets and samples while building their intuitive understanding of how subsets can be more or less representative of their source population. Tally sheets support players in tracking their observations and making evidence-based claims about what their sample reveals about different feature distributions.

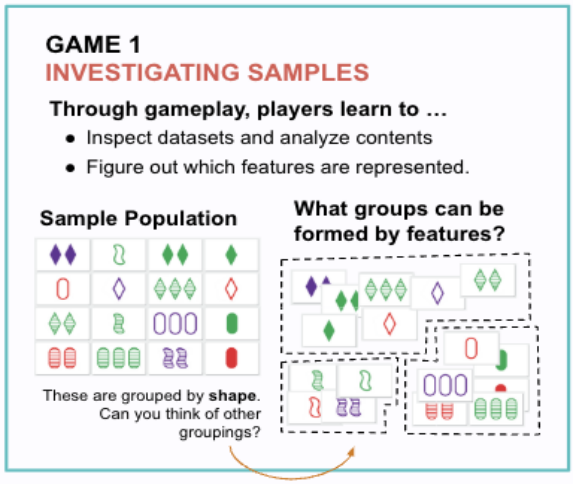

Figure 2: Game 1 learning objectives (top), a visualization of a sample population (left), and an example of how this sample population can be grouped by features.

Known Population & Discoverable Characteristics

An innovation of games is the use of SET cards, which contain all possible combinations of 4 features with 3 values each. This creates a "closed system" where the full population is knowable and discoverable, enabling players to make evidence-based claims. Unlike real-world datasets, where the population characteristics are often unknown or disputed, players can determine what constitutes "representative" sampling using the full deck.

This known population structure allows for discussions about bias and representation grounded in concrete evidence. Players can calculate exact percentages, compare sample distributions to the full population, and test claims about representativeness against objective standards. This empirical foundation fosters understanding difficult to achieve in real-world scenarios where population characteristics are uncertain, and representativeness is contested.

Game 2: Competing Interests

Game 2 introduces strategic decision-making and stakeholder perspectives through role-playing. This is achieved as each player draws an "identity card" from the deck and decides on a goal to have a feature value from that card be either the most or least represented in the sample population as the dataset evolves through gameplay. Players then set the identity card aside for reference throughout the game.

Gameplay progresses as players take turns drawing a new card from the deck and deciding how to play that card based on their identity card and goal. With each turn, a payer has three options: ADD the card to the grid to make the dataset bigger, DISCARD the card in a discard pile face down (removing it from play), or PLACE the card on top of an existing card in the grid to replace another card in the sample dataset. Players make these decisions based on how each action advances their personal objectives.

To simplify gameplay and make the learning objectives more accessible for younger players, the game mechanics allow for players to play several rounds of Game 2. First, all players play to be overrepresented. Then, in a second round, all players may play to be under-represented. This scaffolded approach may make the emergent properties of each goal more apparent to players, regardless of age.

Players observe how the representation of different features changes over time based on collective decision-making, leading to discussions about strategy, unintended consequences, and the challenge of balancing diverse stakeholder needs in dataset construction. The game concludes with reflection on how player goals influenced the final dataset composition and what this reveals about power dynamics in AI development.

Stakeholder Role-Play with Competing Interests

GR requires active manipulation based on stakeholder goals. In game 2, players don't simply analyze existing bias; they create it through their curatorial decisions while pursuing personal objectives that may conflict with others' interests. This creates authentic experiences of how different stakeholders can influence dataset composition in pursuit of their own goals.

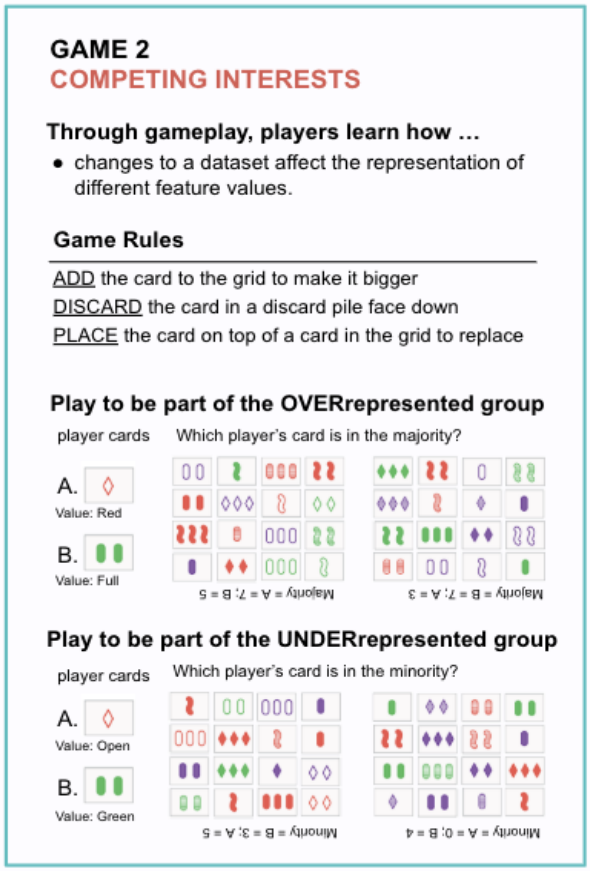

Figure 3: Game 2 learning objectives (top), rules for gameplay (middle top), and an example of how sample populations might look as players play to meet their goals to be over- or under- represented (bottom).

The progression of the games aligns with experiential learning theory (Kolb, 1984; Dewey, 1938), which emphasizes learning through direct experience and reflection on that experience. GR offers reflective discussion upon structured simulations through which players experience phenomena that they cannot easily observe in real-world contexts. GR scaffolds these experiences through a progression aligned with key components of Kolb's experiential learning cycle: concrete experience (manipulating cards), reflective observation (noticing patterns), abstract conceptualization (understanding representation), and active experimentation (exploration and comparison of strategies through game-play and reflective discussion after game-play).

Game 3: Real-World Scenarios

Game 3 applies the dataset manipulation skills from previous games to authentic scenarios where AI systems impact human lives. Using a code sheet (Figure 4) that maps SET card features to human traits (e.g., shape represents height, color represents health status), players reinterpret their manipulated dataset as representing human traits rather than abstract card properties.

Game 3 includes multiple modular scenarios that can be selected based on player age and interests. These scenarios range from more accessible contexts like self-driving car recognition systems (suitable for younger players) to more complex ethical situations such as medical diagnosis tools. We have also considered scenarios involving employment screening algorithms, which may be appropriate for older learners who can engage with nuanced ethical reasoning. Scenarios such as these may become part of future iterations of the games.

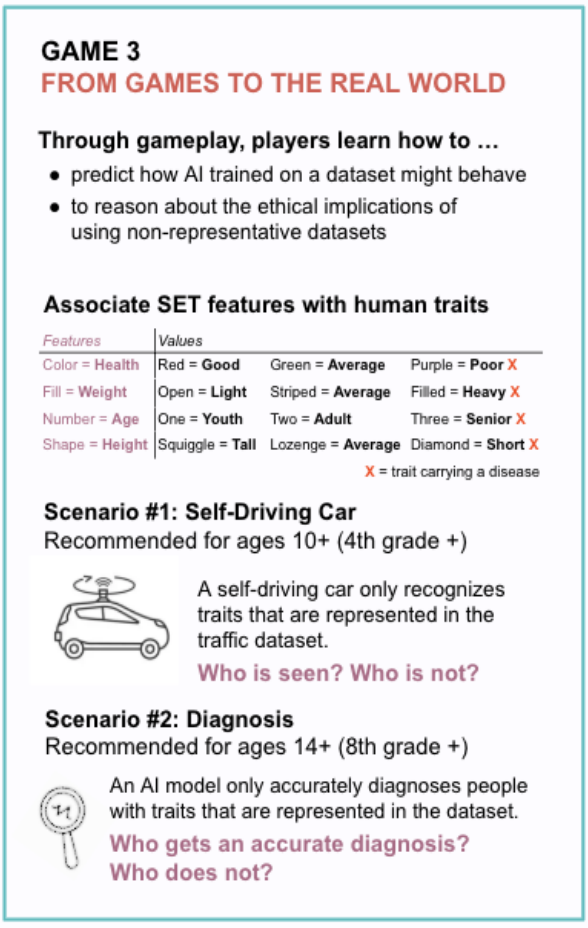

Figure 4: Game 3 learning objectives (top), a code sheet for interpreting card feature values as human traits (middle), and summaries of the two scenarios for gameplay in the context of real-world applications of AI (bottom).

In each scenario, players must decide whether they want their character (identity card) to be represented in the training dataset, considering both the benefits and risks of inclusion. For example, in a medical diagnosis scenario, being represented might enable accurate diagnosis, whereas in scenarios that may be included in future GR game designs involving employment screening, representation might lead to discrimination. Players manipulate the dataset through several rounds of gameplay, making strategic decisions about representation based on the ethical implications of each scenario.

Through guided reflection questions, players explore whose interests are served by different dataset compositions and consider the broader societal implications of their curatorial decisions. This progression across the three games from abstract card manipulation to human-centered ethical reasoning helps players transfer their understanding of representation to real-world AI applications.

The progression of the games aligns with experiential learning theory (Kolb, 1984; Dewey, 1938), which emphasizes learning through direct experience and reflection on that experience. GR offers reflective discussion upon structured simulations through which players experience phenomena that they cannot easily observe in real-world contexts. GR also scaffolds these experiences through a progression aligned with key components of Kolb's experiential learning cycle: concrete experience (manipulating cards), reflective observation (noticing patterns), abstract conceptualization (understanding representation), and active experimentation (e.g., Game 3). The scaffolded structure offers concrete experiences before gameplay engages abstract reasoning.

Progressive Abstraction Through Gameplay

The three-game sequence creates a progression from concrete to abstract thinking. Game 1 focuses on concrete counting and sorting activities, establishing foundational vocabulary and pattern recognition skills. Game 2 introduces strategic manipulation with competing stakeholder interests, adding layers of social complexity and decision-making. Game 3 applies these skills to real-world scenario applications with explicit ethical implications.

Each game builds conceptual complexity while preserving the fundamental card manipulation mechanics, ensuring that cognitive load increases gradually. This progression enables players to develop expertise with the physical materials before engaging with abstract ethical reasoning, creating a stable foundation for increasingly sophisticated discussions about AI.

Concrete-to-Abstract Transfer Mechanism

Game 3 employs an approach that maps abstract SET card features to concrete human traits, enabling players to reinterpret the same dataset through multiple lenses. This makes abstract concepts of dataset bias tangible by allowing players to see how the same data patterns could advantage

or disadvantage different groups depending on the traits being measured.

The mapping system enables sophisticated discussions of proxies and intersectionality while minimizing the risk of overwhelming learners. Players can explore how multiple characteristics interact (tall, young, healthy individuals versus short, elderly, sick individuals) and consider how algorithmic systems might treat people with different combinations of traits differently. This concrete grounding provides essential preparation for understanding real-world dataset bias that operates through interactions of multiple variables; thus, GR addresses *Core Component #1* by engaging players in the examination of a simplified AI technology ecosystem.

Scaffolded Ethical Reasoning

The modular scenarios in Game 3 creates a progression from ethically neutral contexts (self-driving cars that need to recognize all pedestrians) to personally relevant situations (medical diagnosis where inclusion enables treatment), and in future versions of the games, potentially harmful applications (employment screening that could enable discrimination). This scaffolding allows players to develop ethical reasoning skills gradually, building confidence with straightforward scenarios before tackling more complicated moral dilemmas.

Built-in discussion prompts guide players through structured ethical reasoning processes, encouraging perspective-taking and consideration of voices not present at the table. The scenarios prompt players to think beyond their immediate experience to consider how their dataset manipulation decisions might impact absent stakeholders, fostering the kind of inclusive ethical reasoning essential for responsible AI development; thus, GR addresses *Core Component #6,* through conversations that prompt players to consider how to minimize bias and the risk of harm caused by AI.

Together, all three games incorporate several established pedagogical design principles that enhance their potential for educational effectiveness and practical implementation. The games' multi-modal approach engages different learning preferences through kinesthetic manipulation, visual pattern recognition, and collaborative discussion, providing multiple pathways for learners to access and understand complex AI concepts. Flexible implementation design (e.g., the modularity of the games, the opportunities for scaffolded gameplay) ensures broad applicability across varied educational settings and age groups, while built-in measurement tools (i.e., the tally sheet that allows players to quantify the shifting sample population distribution) make learning visible. These core design features establish a solid pedagogical foundation that supports the games' more innovative elements (described below) creating an educational resource that is both theoretically grounded and practically viable for widespread adoption across diverse learning environments. Below we detail the games' more innovative elements.

Discussion

While GR remains in active development, the initial design features suggest that the games are well-positioned to contribute to the field of AI education in several ways.

Exploring the Inner Workings of AI Models

The games answer the call for tools designed for novices to build and explore how models work (Gresse von Wangenheim et. al. 2021) and addresses the need for better educational tools to support learners in not just using AI but creating AI models (Carney et. al. 2020; Morales-Navarro et. al. 2025). GR achieves this by positioning learners as active dataset curators who directly influence model training data rather than passive consumers of pre-built AI systems.

GR is not alone in answering this call. The Tensorflow Neural Network Playground, Google's Teachable Machine, and the interactive visualizations available through Google's Machine Learning course, for example, all offer demonstrations of the inner workings of AI models, specifically the training process. Platforms such as these allow learners to explore the functionality of fundamental algorithms in AI and machine learning technologies. Yet, they also obscure the dataset construction process behind user-friendly interfaces. For example, with Google's Teachable Machine, users can upload data to train models, but this is a simulation using transfer learning; users cannot explore the transformer model nor the MobileNet training dataset. In these ways, interactive simulations rarely afford learners opportunities to examine the representativeness of their data samples.

GR complements rather than competes with these digital exploration tools. By making dataset construction the primary focus, GR serves as conceptual preparation for more technical experiences, helping learners approach digital AI tools with greater critical awareness of the data they are using.

Practical Application #1 – Educators, families and caregivers may consider using GR as foundational preparation before introducing digital model-building tools, enabling learners to understand the human decisions embedded in training datasets before engaging with the technical aspects of AI model training.

Hands-On, Low-Tech AI Education

Many existing AI education tools require some technical expertise, software environments, device access, internet connectivity, and occasional technical troubleshooting. GR eliminates these technological barriers while introducing foundational concepts essential to AI model development. The need for this type of intervention is well established in

computer science education (Linder et. al. 2019) and more recently in AI education (Dai 2025; Lim et. al. 2025). GR joins a well-developed collection of unplugged activities designed to advance understanding of AI concepts, i.e., facial recognition (Lim et. al. 2024), decision trees and reinforcement learning (Lindner et. al. 2019), and semantic networks and knowledge representations (Long et. al. 2021).

Practical Application #2 – Educators, families, and caretakers might choose low-tech approaches like GR when working in resource-constrained environments and when introducing AI concepts to younger learners, who may not yet be ready for programming interfaces.

Sociotechnical AI Education

GR emphasizes the social contexts and human decisions that shape how algorithms operate in the real world. The games center the social nature of data curation (Tseng et. al. 2024; Hardy et. al. 2020) inherent in the design of many computational systems. GR accomplishes this through explicit stakeholder role-play that simulates real-world tensions in AI development. The corresponding guide facilitates group discussion and problem-solving around ethical dilemmas, moving beyond technical understanding to engage critical examination of competing interests in AI.

Practical Application #3 – Educators, families and caregivers might choose GR over other collaborative AI activities when seeking to complement technical AI education with critical social analysis, when working with diverse groups where different lived experiences can enrich discussions of bias and representation, or when addressing real-world AI controversies where understanding multiple stakeholder perspectives can foster informed civic participation.

Bringing AI Education into Informal Settings

Many AI education resources target formal classroom settings with teacher-led instruction, leaving families, caregivers and community organizations without accessible tools for engaging these critical topics. AI systems increasingly impact family and community life through applications in healthcare, employment, criminal justice, and social services, yet community stakeholders are too often excluded from conversations about AI development.

There is some empirically grounded guidance for parents and families as to how to have conversations about AI with children and youth (e.g., Rubin et. al. 2025). What's more, there are some resources, like the Day of AI and Common-Sense Media Family Guide, that facilitate hands-on activities, collaborative exploration of AI's capabilities, and discussion starters for families to engage in conversation about AI and how it works with youth.

GR joins this effort and goes a step further by creating a shared, scaffolded exploration driven experience through which caregivers and adolescents can authentically discover the limitations of AI and practice making decisions together that are immediately relevant to their daily lives (e.g., Do you want your data included in this dataset?).

Practical Application #4 – Communities and caregivers might choose GR when seeking to build collective understanding of AI impacts on their daily lives, when addressing specific concerns about algorithmic bias in local systems, or when working to preserve cultural knowledge and data sovereignty in the face of expanding AI applications.

Future Directions

GR development is ready to follow several promising avenues. Pilots and efficacy studies across diverse populations and educational contexts are needed to demonstrate impact on players' knowledge and engagement, specifically regarding understanding of AI concepts such as bias and representation in AI training datasets. The core game mechanics could be extended to teach other AI concepts beyond dataset representation, such as algorithmic fairness, model interpretability, or AI system lifecycle management. Finally, systematic adaptation of the games for different age ranges and cultural contexts could broaden their accessibility and relevance, ensuring that this approach to AI ethics education serves diverse learners across varied educational settings.

Conclusion

GR and the corresponding guide contribute to the growing field of K-12 AI ethics education by addressing critical gaps in current AI education resources. Through hands-on gameplay using the ready-made SET card deck and simple game mechanics, players learn to understand abstract AI concepts through manipulation of datasets, transforming complicated ideas about representation and bias into tangible, playable experiences.

As AI systems become increasingly prevalent in adolescents' daily lives, GR offers educators, families, and caregivers a powerful resource for hosting increasingly important conversations with adolescents about data privacy as well as opportunities to develop critical thinking skills necessary to navigate an AI-influenced world. GR is designed to make ethical concepts accessible to young learners, contributing meaningfully to the essential work of preparing adolescents for responsible engagement with AI.

Acknowledgments

This work would not have been possible without the generous support of the Kapor Foundation. Special thanks to Maria Levy for the visualization and graphic design of the Games of Representation guide.

References

Bell, T.; Alexander, J.; Freeman, I.; and Grimley, M. 2009. Computer science unplugged: School students doing real computing without computers. *New Zealand Journal of Applied Computing and Information Technology* 13(1): 20-29.

Carney, M.; Webster, B.; Alvarado, I.; Phillips, K.; Howell, N.; Griffith, J.; Jongejan, J.; Pitaru, A.; and Chen, A. 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, https://doi. org/10.1145/3334480.3382839.

Dai, Y. 2025. Integrating unplugged and plugged activities for holistic AI education: An embodied constructionist pedagogical approach. *Education and Information Technologies* 30: 6741–6764. https://doi.org/10.1007/s10639-024-13043-w.

Day of AI. 2025. Day of AI and Common Sense Media Family Guide. https://dayofai.org/families/day-of-ai-common-sense-media-ai-literacy-toolkit-for-families/

Dewey, J. 1938. *Experience and education*. Kappa Delta Pi.

Google. 2017. Teachable Machine. https://teachablemachine.withgoogle.com/.

Google. 2024. Machine Learning Crash Course: Neural Networks. https://developers.google.com/machine-learning/crash-course/neural-networks/nodes-hidden-layers.

Gresse von Wangenheim, C.; Hauck, J. C.; Pacheco, F. S.; and Bertonceli Bueno, M. F. 2021. Visual tools for teaching machine learning in K-12: A ten-year systematic mapping. *Education and Information Technologies 26*(5): 5733-5778. https://doi.org/10.1007/ s10639-021-10570-8.

Hardy, L.; Dixon, C.; and Hsi, S. 2020. From data collectors to data producers: Shifting students' relationship to data. *Journal of the Learning Sciences* 29(1): 104-126. https://doi.org/10.1080/ 10508406.2019.1678164.

Kolb, D. A. 1984. *Experiential learning: Experience as the source of learning and development*. Prentice Hall.

Lee, I.; Ali, S.; Zhang, H.; DiPaola, D.; and Breazeal, C. 2021. Developing middle school students' AI literacy. In Proceedings of the 52nd ACM technical symposium on computer science education: 191-197. https://doi.org/10.1145/3408877.3432513

Lim, H.; Min, W.; Vandenberg, J.; Cateté, V.; and Mott, B. 2024. Unplugged K-12 AI Learning: Exploring Representation and Reasoning with a Facial Recognition Game. In Proceedings of the AAAI Conference on Artificial Intelligence 38(21): 23285-23293. https://doi.org/10.1609/aaai.v38i21.30376.

Lim, H.; Boulden, D.; Vandenberg, J.; Cateté, V.; Min, W.; and Mott, B. 2025. Shaping AI Interest in Rural Middle Schools with Unplugged Learning: Gender Differences and Teacher Insights. In Proceedings of the AAAI Conference on Artificial Intelligence 39(28): 29153-29161. https://doi.org/10.1609/aaai.v39i28.35188.

Lindner, A.; Seegerer, S.; and Romeike, R. 2019. Unplugged Activities in the Context of AI. In Informatics in Schools. New Ideas in School Informatics. Lecture Notes in Computer Science, edited by S. Pozdniakov and V. Dagienė, V. 11913, 123-135. https://doi.org/10.1007/978-3-030-33759-9_10

Long, D.; Blunt, T.; and Magerko, B. 2021. Co-designing AI literacy exhibits for informal learning spaces. In Proceedings of the ACM on Human-Computer Interaction, 5(CSCW2), 1-35. https://doi.org/10.1145/3476034.

Morales-Navarro, L.; Kafai, Y. B.; Yang, E.; and Suryana, A. 2025. What Can Youth Learn About Artificial Intelligence and Machine Learning in One Hour? Examining How Hour of Code Activities Address the Five Big Ideas of AI. In Proceedings of the AAAI Conference on Artificial Intelligence, 39(28), 29195-29202. https://doi.org/10.1609/aaai.v39i28.35193

Rubin, J. D.; Lombard, E. J.; Chen, K.; Divanji, R.; Min, A.; Dotch, E.; Han, A.; Dickerson, K.; Hayes, G.; and Odgers, C. 2025. Navigating AI as a family: Caregivers' perspectives and strategies [White paper]. https://www.foundry10.org/resources/navigating-ai-as-a-family.

Sysoev, I.; Gray, J. H.; Fine, S.; Makini, S. P.; and Roy, D. 2022. Child-driven, machine-guided: Automatic scaffolding of constructionist-inspired early literacy play. *Computers & Education*, 182, 104434. https://doi.org/10.1016/j.compedu.2022.104434.

Tensorflow. 2022. Tensorflow Neural Network Playground. https://github.com/tensorflow/playground.

Tseng, T.; Davidson, M. J.; Morales-Navarro, L.; Chen, J. K.; Delaney, V.; Leibowitz, M.; Beason, J; and Shapiro, R. B. 2024. Co-ML: Collaborative machine learning model building for developing dataset design practices. *ACM Transactions on Computing Education*, *24*(2), 1-37. https://doi.org/10.1145/3641552.

Zhang, H.; Lee, I.; Ali, S.; DiPaola, D.; Cheng, Y.; and Breazeal, C. 2023. Integrating ethics and career futures with technical learning to promote AI literacy for middle school students: An exploratory study. International Journal of Artificial Intelligence in Education, 33(2), 290-324. https://doi.org/10.1007/s40593-022-00293-3

Zhang, H.; Lee, I.; and Moore, K. 2024. An effectiveness study of teacher-led AI literacy curriculum in K-12 classrooms. In Proceedings of the AAAI conference on artificial intelligence, 38(21): 23318-23325. https://doi.org/10.1609/aaai.v38i21.30380