Emerging Uses of AI-Generated Images for Equitable and Transparent Simulations

Philippe J. Giabbanelli[1], Kourosh Shoele[2], Megan A. Witherow[1]
[1]Virginia Modeling, Analysis, and Simulation Center (VMASC), Old Dominion University
1030 University Blvd, Suffolk, VA 23435, USA
[2]Department of Mechanical Engineering, Florida State University, Tallahassee, Florida, USA
pgiabban@odu.edu, ,mwithero@odu.edu, kshoele@fsu.edu

Abstract

Despite the maturity of AI image generation, integration of generative AI in M&S has primarily been limited to text. This paper presents a vision for the use of AI-generated images in M&S with an emphasis on equity and transparency. We suggest several emerging use cases including AI-generated images acting as interfaces between agent-based models and physics-based simulations, encouraging empathetic decision-making by visualizing individual agents, and promoting transparency with symbolic representations that complement textual descriptions of abstract model processes. Finally, we discuss the mitigation of ethical issues related to the deployment of AI-generated images in M&S.

Introduction

Commissioners, researchers, and stakeholders in Modeling & Simulation (M&S) are actively integrating generative AI tools throughout stages of the M&S process, as illustrated in Figure 1. Integration has primarily taken the form of text inputs to large language models such as OpenAI's GPT to produce text outputs. For example, a model may be created by converting a text narrative (i.e., input) into a structured form (e.g., causal or cognitive map, process models), which is output as a text file (e.g., a bulleted list of edges) rather than as an image (node link diagram).

In contrast, there is currently limited integration of AI-driven text-to-image generative models (a.k.a. 'AI-based image generation') in the M&S process. This is not due to a lack of maturity in AI technology for image generation or lack of familiarity with such possibilities for practitioners. AI-generated images are now commonplace. They are active topics of debate when it comes to deepfakes (Lao, Hirvonen, and Larsson 2025; Grossman and Grimm 2025) and AI-generated images are routinely used in the entertainment industry (Lapointe et al. 2025; Sew et al. 2025; Zhang et al. 2025). Manuscripts have used AI-generated images either covertly (Dash et al. 2024; Gu et al. 2022; Skulmowski and Engel-Hermann 2025) or for scientific needs, such as augmenting the volume and resolution of training data for image processing algorithms (Zhou et al. 2025; Xie et al. 2025). In the M&S process, images are mostly encountered as scientific visualizations to examine simulation outcomes. For instance, we can observe aggregated spatio-temporal patterns as simulated individuals move through space in an Agent-Based Model (Grignard and Drogoul 2017) or we can see how a forest fire spreads under different conditions in a cellular automaton (Giabbanelli and Baniukiewicz 2019). These images may serve as input to AI tools, e.g., to check the validity of a simulation (Wozniak and Giabbanelli 2021) or to learn about simulations (Flandre and Giabbanelli 2024), but are they are *produced by the simulation* rather than generative AI tools. We thus posit that the lack of integration of AI-based image generation with the M&S process may stem from an insufficient awareness of potential uses.

In this vision paper, we present several ways in which AI-based image generation can support M&S. Specifically, we examine how this integration can improve the equity of simulation-based policy recommendations and the ability of users to engage with simulations, which is important to establish trust and promote participation.

AI Images as Interfaces for Physics-Based Simulations and Agent-Based Modeling

Agent-Based Modeling is a M&S technique in which individual entities (virtual humans in our case) interact with each other and with their physical surroundings. This technique is commonly applied in socio-environmental problems and recently came to attention outside of the simulation community due to its widespread use in computational epidemiology for COVID-19 (Lorig et al. 2021; Belfrage et al. 2024). In agent-based models (ABMs), agents are defined by features such as age category or sex, and they make decisions (e.g., whether to wear a face mask) based on the information available to them, such as observations of their peers' behaviors and adherence to policies. While ABMs *rely* on data for calibration and validation (Badham et al. 2018), AI has long been used in this field to make ABMs more data-driven (Kavak et al. 2018) by automating model implementation (Martínez et al. 2024), replacing the agents' decision-making modules with machine learning models (Beerman et al. 2023), or building a more computationally efficient version of the simulation via a machine learning surrogate (Fonseca et al. 2025). Despite the many ways in which generative AI has been used with ABMs (Vanhée et al. 2025), AI-generated images are noticeably absent.

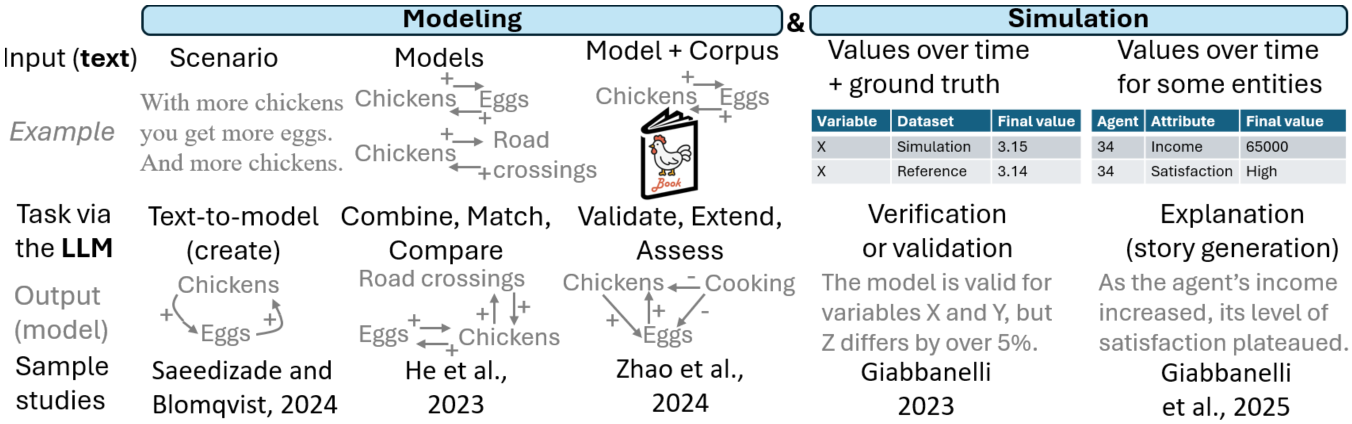

Figure 1: Multiple studies have automated tasks of the Modeling & Simulation process by using LLMs. In all of these instances, input and output are both text files, which can be unstructured (e.g., a scenario) or encode a graph (e.g., JSON, CSV).

Generating an image for each agent is straightforward using existing open-source technologies (Fig. 2, top), as the characteristics of an agent can be assembled as a list and passed to a text-to-image tool such as Stable Diffusion to generate photorealistic portraits. The main question is: *what should we do with the images*? Using simulations of respiratory diseases (e.g., COVID-19) as an example, the image of an agent can be transformed into a 3D morphing face model to simulate mask deployment (Fig. 2, bottom) and the protection that it confers to the agent (inhaling and exhaling viral particles) (Solano, Mittal, and Shoele 2021; Wang, Solano, and Shoele 2022; Anand and Shoele 2025). Physical differences associated with demographic traits (e.g., age, race, and sex) affect the fit of a mask and its efficacy. For instance, as people age, subcutaneous fat redistributes, altering facial contours around the chin and neck (potentially reducing mask leakage by filling gaps), while age-related asymmetry can cause uneven mask contact pressure (Griffin et al. 2024; Solano and Shoele 2022). Using AI-based image generation would thus provide estimates of mask efficiency across diverse agents, which allows us to *identify who benefits from interventions and who remains at risk*. Such simulations may reveal that some 'universal' policies are not universally protective (i.e., unequal outcomes among groups), which may suggest opportunities for policy equity, such as providing high-quality masks to groups with higher leakage rates.

Through this example, we emphasize that *AI-generated images may play an essential role in connecting separate strands of research*. On the one hand, the infection in ABMs is only based on proximity (Schroeder et al. 2022; Ebrahimi et al. 2024; Reveil and Chen 2022) and aerosol diffusion in the environment (Xue et al. 2024), thus ignoring differences between the agents and lacking support for equitable simulations. On the other hand, computational fluid dynamics can perform simulation of mask fits using well-known facial image databases such as the American Multiracial Face Database (Chen, Norman, and Nam 2021), the Basel Face Database (Paysan et al. 2009), or the Chicago Face Database (Ma, Correll, and Wittenbrink 2015), but it lacks the ABM component to simulate viral spread among agents or how the use of masks depends on places (e.g., in a regulated public hub vs. at home). In sum, AI images are the piece of the puzzle that connects different types of simulation to arrive at a hybrid solution that assesses equity.

Bridging Model Abstractions and Human Realities Through AI-Generated Images

Simulation models have been used to study the impact of potential decisions across multiple aspects of equity, including social position, human capital, socioeconomic and political context (Mui et al. 2025; Giabbanelli 2024). In particular, ABMs provide decision-makers with tools to estimate the potential effect of different resource allocations across social groups (Tomasiello, Giannotti, and Feitosa 2020). However, these tools could create a disconnect between decision-makers and the individuals whose lives are impacted by decisions, leading to two specific problems. First, decision-makers may account for only a subset of variables at play in an individual's life, since a model is necessarily a simplification of reality. For instance, a simulation can forecast that a significant reduction in the evacuation budget may still suffice to move people away from danger, but it would not capture the struggle of families who face transportation barriers, require regular access to medication, or fear interacting with authorities due to immigration status. The aggregate metrics of a simulation may hide the vulnerability of marginalized groups, leading decision-makers to rely on optimistic projections. Second, neither the tools (which were designed by people) nor their users are free of biases, but models may be wrongly used to deflect responsibility (e.g., "the model supports this choice"). For example, a decision-maker may handle trade-offs in favor of a group to which they feel more connected, e.g. leading to gentrification.

Generative AI has been used to address the lack of transparency or promote empathy in decision making. For example, several studies have examined how LLMs 'feel' (Sorin et al. 2024; Huang et al. 2024) in reaction to a narrative or whether they can craft stories that elicit certain feelings in

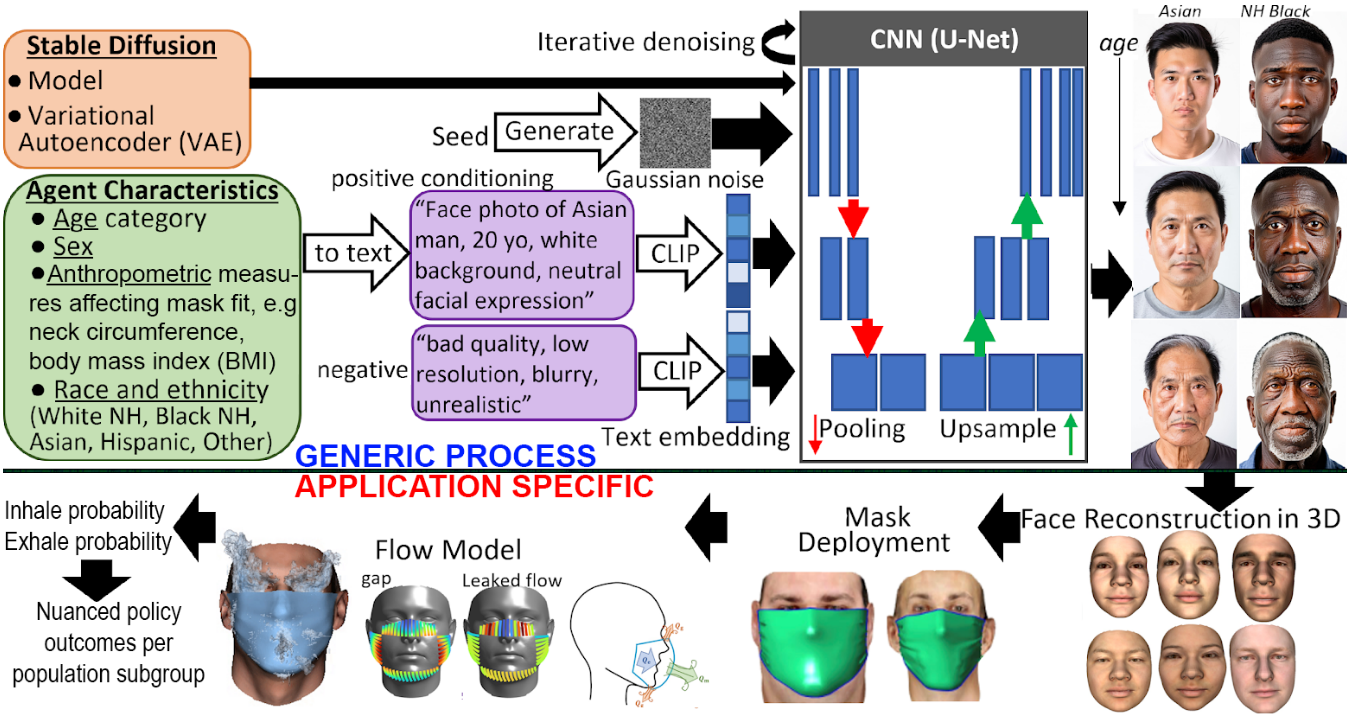

Figure 2: The general process to turn an agent (as described by features such as age and sex) into an image is to gather the features' values as text and provide them to a text-to-image model (e.g., Stable Diffusion, Flux, DALL-E, Midjourney). The image can be used for downstream applications, which can involve other simulation paradigms such as fluid dynamics.

human readers (Shen et al. 2024). Our recent work on ABMs translates the simulated journey of an agent during a hurricane evacuation into a complete narrative, which was evaluated by human readers as showing genuine emotions (Giabbanelli et al. 2025). Several works have also been devoted to improving transparency in M&S by using LLMs to explain the structure of the model (Tel and Minor 2025), which has quickly improved from early report generation focusing on sentences (i.e., one giant paragraph) to current systems that create executive summaries (Gandee and Giabbanelli 2024).

Although the above solutions approached the lack of transparency and the need for empathy *solely through text*, we posit that AI-generated images can benefit both cases.

Photorealistic images or videos may *evoke empathy* more quickly and robustly than text alone. For example, instead of reading several pages about the journey of each agent, decision-makers could *see* key moments in an agent's life in a fraction of the time. The main research challenges are not about the technical feasibility, since models such as Stable Diffusion or Flux can provide photorealism. Rather, deployment presents unique challenges because we cannot use *any* visual to narrate a simulated agent's journey and we should not ignore the important matter of *attribution*. First, while visuals may benefit information recall and mobilize several cognitive subsystems, these benefits can be lost if the visuals are overly complex, leading to cognitive overload. Studies are thus needed to evaluate how the design of images would impact cognitive load, empathy, and ultimately the decisions made. Second, Dorigoni and Giardino (2025) point out a key challenge: when AI-generated images include anthropomorphic features, viewers may perceive them as inauthentic once they realize they were created by AI. To address this, the authors recommend "to pair it with human oversight, disclose authorship clearly, or frame the message as co-authored with a trusted human". While analyses of collaborative and participatory modeling show that people are involved mainly at the early ideation stage (Manellanga and David 2024), we see the need for human involvement as an opportunity for engagement throughout the M&S process, culminating in a co-design approach to narrating results via AI-generated images that accurately mirror lived experiences. This human oversight in AI generation would also allow communities to contribute rich personal insights that could not be captured through a simulation alone, whether due to limited data or necessary simplifications in the model.

Promoting transparency in simulation design (e.g., the rules that underpin the agents' behavior) requires a different tool set and calls for different research questions than conveying empathetic moments in the simulated journeys of individual agents. Photorealistic images excel at creating emotional immediacy as they show the consequences of concrete events experienced by the agents. In contrast, the rules of a model are more abstract, such as 'if most peers express a social norm that is different from mine then I will gradually move in their direction'. These abstractions call for symbolic representations (e.g., using tools such as DALL·E). While we envision that images or videos of agents may be provided *instead* of text (for empathy), we posit that representa-

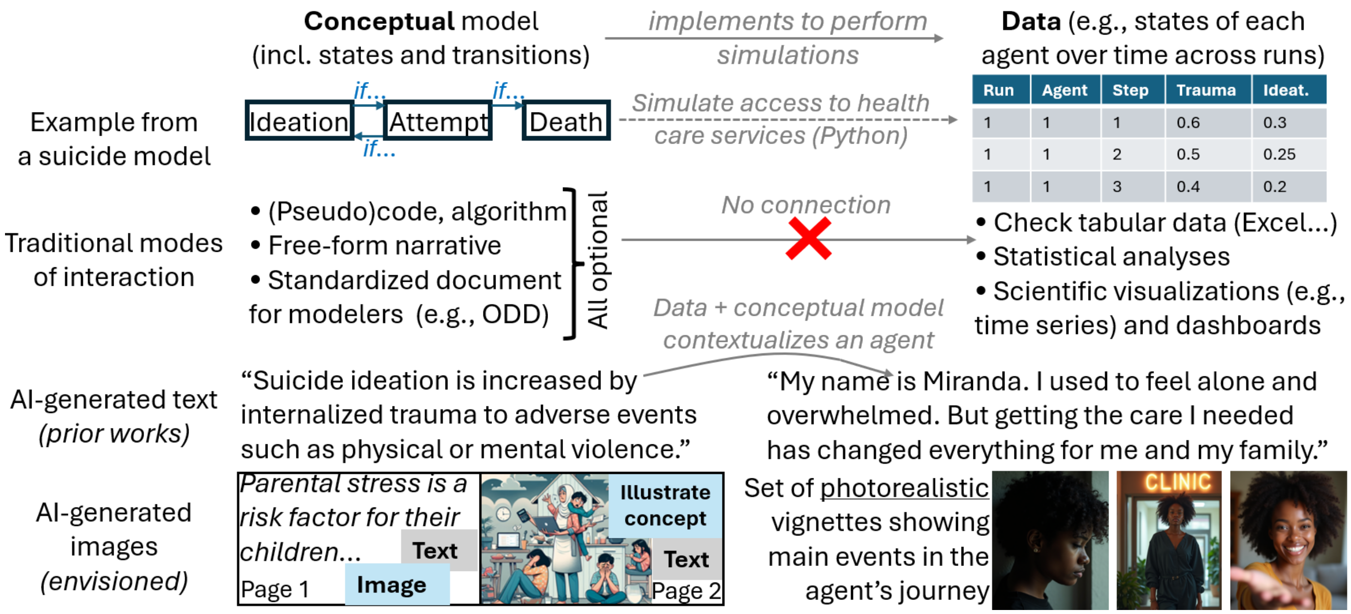

Figure 3: Various forms of documentation, spreadsheets, statistics, and scientific visualizations may provide insight into a model's rules or its outputs. However, these formats may not be directly usable by a decision-maker. Prior works have used text summaries generated by large language models to provide more accessible, but often lengthy descriptions. We suggest the complementary use of symbolic representations (e.g., from DALL-E) and text to explain how the model works, and the use of photorealistic images (e.g., from Flux) to promote immersion in an agent's simulated journey.

tions of abstract processes may best promote understanding and analysis when they *complement* a textual description, as portrayed in Figure 3. So far, the use of AI-generated data to promote understanding and analysis has primarily been through data visualization, which is concerned with facilitating insights by creating the right charts (Wang, Sundstedt, and Garro 2025). The generation of images about the *process* of a model rather than the *results* of a simulation raises specific, interrelated research questions: *which parts* of a model should be reinforced through visuals, and how to *integrate* visuals with the text? Early integration experiments suggest that not every part of a model is a good candidate for visuals (e.g., some are too metaphorical and do not contribute to understanding the model) or even feasible (e.g., a simulation model of suicide may trigger banned keywords in an AI image generation tool) (Gandee et al. 2024).

Mitigating Ethical Issues in Deploying AI-Generated Images for Simulations

The emergence of text-to-image models has sparked discussions on ethical and responsible use of generated images (Hagendorff 2024), many of which are pertinent for their use in M&S. AI tools generate new images by learning the internal representation of a distribution containing augmented and *real images*. Biases associated with the quantity and distribution of training data representing different racial, ethnic, age, sex, and other demographic groups have been documented across multiple AI text-to-image tools (Luccioni et al. 2023). Factors such as feature diversity (e.g., facial features), data quality (e.g., image resolution), and sample selection (e.g., data sources) may further affect the integrity of AI output, especially considering data dependencies with demographic groups (e.g., less variance of facial features in images of minority groups) (AlDahoul, Rahwan, and Zaki 2025). Furthermore, generative AI tools are prone to 'hallucinations', which may help to represent new situations but can decrease immersion when producing unusual bodies (e.g., hands with three fingers) or improbable interactions with the surroundings (e.g., levitating objects) (Aithal et al. 2024). Regions of the real-world data distribution that are poorly sampled for training AI tools may be more susceptible to privacy concerns such as extraction of outliers or duplicated training samples (e.g., to increase representation of underrepresented groups) (Carlini et al. 2023), and security concerns such as adversarial attacks for targeted manipulation of generated content (Zhuang, Zhang, and Liu 2023).

If not carefully incorporated into the M&S process, biases in the sampling of training data may lead to less accurate or representative simulation results for underrepresented groups (i.e., when AI images act as an interface to physics-based modeling); reinforce stereotypes, ideological leanings, and marginalization of underrepresented communities (i.e., visualization of an agent's journey); or introduce new privacy and security risks for marginalized populations. Thorough bias characterization and de-biasing, hallucination detection, privacy and security safeguards, and communication of risks and benefits will help ensure safe, equitable, and transparent use of AI-generated images in M&S.

References

Aithal, S. K.; Maini, P.; Lipton, Z. C.; and Kolter, J. Z. 2024. Understanding Hallucinations in Diffusion Models through Mode Interpolation. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 134614–134644. Curran Associates, Inc.

AlDahoul, N.; Rahwan, T.; and Zaki, Y. 2025. AI-generated faces influence gender stereotypes and racial homogenization. *Scientific Reports*, 15(1): 14449.

Anand, A.; and Shoele, K. 2025. Identifying the effectiveness of face mask in a large population with a network-based fluid model. *PLoS One*, 20(6): e0324229.

Badham, J.; Chattoe-Brown, E.; Gilbert, N.; Chalabi, Z.; Kee, F.; and Hunter, R. F. 2018. Developing agent-based models of complex health behaviour. *Health & place*, 54: 170.

Beerman, J. T.; et al. 2023. A framework for the comparison of errors in agent-based models using machine learning. *Journal of computational science*, 72: 102119.

Belfrage, M.; et al. 2024. Simulating change: A systematic literature review of agent-based models for policy-making. In *2024 Annual Modeling and Simulation Conference (ANNSIM)*, 1–13. IEEE.

Carlini, N.; Hayes, J.; Nasr, M.; Jagielski, M.; Sehwag, V.; Tramèr, F.; Balle, B.; Ippolito, D.; and Wallace, E. 2023. Extracting Training Data from Diffusion Models. In *32nd USENIX Security Symposium*, 5253–5270. Anaheim, CA.

Chen, J. M.; Norman, J. B.; and Nam, Y. 2021. Broadening the stimulus set: Introducing the American multiracial faces database. *Behavior Research Methods*, 53(1): 371–389.

Dash, K. S.; et al. 2024. We are entering a new era of problems: AI-generated images in research manuscripts.

Dorigoni, A.; and Giardino, P. L. 2025. The illusion of empathy: evaluating AI-generated outputs in moments that matter. *Frontiers in Psychology*, 16: 1568911.

Ebrahimi, A. H.; Alesheikh, A. A.; Hooshangi, N.; Sharif, M.; and Mollalo, A. 2024. Modeling COVID-19 Transmission in Closed Indoor Settings: An Agent-Based Approach with Comprehensive Sensitivity Analysis. *Information*, 15(6): 362.

Flandre, N. Y.; and Giabbanelli, P. J. 2024. Can large language models learn conceptual modeling by looking at slide decks and pass graduate examinations? an empirical study. In *Int. Conf. on Conceptual Modeling*, 198–208. Springer.

Fonseca, L. L.; Böttcher, L.; Mehrad, B.; and Laubenbacher, R. C. 2025. Optimal control of agent-based models via surrogate modeling. *PLoS Computational Biology*, 21(1): e1012138.

Gandee, T. J.; and Giabbanelli, P. J. 2024. Combining natural language generation and graph algorithms to explain causal maps through meaningful paragraphs. In *Int. Conf. on Conceptual Modeling*, 359–376. Springer.

Gandee, T. J.; et al. 2024. A Visual Analytics Environment for Navigating Large Conceptual Models by Leveraging Generative Artificial Intelligence. *Mathematics*, 12(13): 1946.

Giabbanelli, P. J. 2023. GPT-based models meet simulation: How to efficiently use large-scale pre-trained language models across simulation tasks. In *2023 Winter simulation conference (WSC)*, 2920–2931. IEEE.

Giabbanelli, P. J. 2024. Emerging directions in leveraging machine intelligence for explainable and equity-focused simulation models of mental health. In *Proc. AAAI Symposium Series*, volume 4, 298–302.

Giabbanelli, P. J.; and Baniukiewicz, M. 2019. Visual analytics to identify temporal patterns and variability in simulations from cellular automata. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 29(1): 1–26.

Giabbanelli, P. J.; Daumas, C.; Flandre, N. Y.; Pitkar, A.; and Vazquez-Estrada, J. 2025. Promoting empathy in decision-making by turning agent-based models into stories using large-language models. *Journal of Simulation*.

Griffin, J. S.; McInroe, E. M.; Pennington, E. R.; Steinhardt, W.; Chen, H.; Prince, S. E.; and Samet, J. M. 2024. Craniometric determinants of the fitted filtration efficiency of disposable masks. *Frontiers in Public Health*, 12: 1444411.

Grignard, A.; and Drogoul, A. 2017. Agent-Based Visualization: A Real-Time Visualization Tool Applied Both to Data and Simulation Outputs. In *AAAI Workshops*.

Grossman, M.; and Grimm, P. 2025. Judicial Approaches to Acknowledged and Unacknowledged AI-Generated Evidence. *Science and Technology Law Review*, 26(2).

Gu, J.; Wang, X.; Li, C.; Zhao, J.; Fu, W.; Liang, G.; and Qiu, J. 2022. AI-enabled image fraud in scientific publications. *Patterns*, 3(7).

Hagendorff, T. 2024. Mapping the Ethics of Generative AI: A Comprehensive Scoping Review. *Minds and Machines*, 34(4): 39.

He, Y.; Chen, J.; Dong, H.; and Horrocks, I. 2023. Exploring large language models for ontology alignment. *arXiv:2309.07172*.

Huang, J.-t.; Lam, M. H.; Li, E. J.; Ren, S.; Wang, W.; Jiao, W.; Tu, Z.; and Lyu, M. R. 2024. Apathetic or empathetic? evaluating llms' emotional alignments with humans. *Advances in Neural Information Processing Systems*, 37: 97053–97087.

Kavak, H.; Padilla, J. J.; Lynch, C. J.; and Diallo, S. Y. 2018. Big data, agents, and machine learning: towards a data-driven agent-based modeling approach. In *Proc. annual simulation symposium*, 1–12.

Lao, Y.; Hirvonen, N.; and Larsson, S. 2025. Everyday encounters with deepfakes: young people's media and information literacy practices with AI-generated media. *Journal of Documentation*, 81(7): 216–235.

Lapointe, V. A.; Dubé, S.; Rukhlyadyev, S.; Kessai, T.; and Lafortune, D. 2025. The present and future of adult entertainment: A content analysis of AI-generated pornography websites. *Archives of Sexual Behavior*, 1–19.

Lorig, F.; et al. 2021. Agent-based social simulation of the COVID-19 pandemic: A systematic review. *JASSS: Journal of Artificial Societies and Social Simulation*, 24(3).

Luccioni, S.; Akiki, C.; Mitchell, M.; and Jernite, Y. 2023. Stable Bias: Evaluating Societal Representations in Diffusion Models. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 56338–56351. Curran Associates, Inc.

Ma, D. S.; Correll, J.; and Wittenbrink, B. 2015. The Chicago face database: A free stimulus set of faces and norming data. *Behavior research methods*, 47(4): 1122–1135.

Manellanga, R.; and David, I. 2024. Participatory and collaborative modeling of sustainable systems: A systematic review. In *Proc. ACM/IEEE 27th Int. Conf. on model driven engineering languages and systems*, 645–654.

Martínez, J.; Llinas, B.; Botello, J. G.; Padilla, J. J.; and Frydenlund, E. 2024. Enhancing GPT-3.5's Proficiency in Netlogo Through Few-Shot Prompting and Retrieval-Augmented Generation. In *2024 Winter Simulation Conference (WSC)*, 666–677. IEEE.

Mui, Y.; Winkler, M. R.; Hunt, S. L.; Gittelsohn, J.; and Tracy, M. 2025. Simulated retail food environments: A literature review of systems science approaches to advance equity in access to healthy diets. *Obesity Reviews*, 26(5): e13887.

Paysan, P.; Knothe, R.; Amberg, B.; Romdhani, S.; and Vetter, T. 2009. A 3D face model for pose and illumination invariant face recognition. In *6th IEEE Int. Conf. on advanced video and signal based surveillance*, 296–301.

Reveil, M.; and Chen, Y.-H. 2022. Predicting and preventing COVID-19 outbreaks in indoor environments: An agent-based modeling study. *Scientific Reports*, 12(1): 16076.

Saeedizade, M. J.; and Blomqvist, E. 2024. Navigating Ontology Development with Large Language Models. In *European Semantic Web Conf.*, 143–161. Springer.

Schroeder, S. A.; Vendome, C.; Giabbanelli, P. J.; and Montfort, A. M. 2022. Towards reusable building blocks to develop COVID-19 simulation models. In *2022 Winter Simulation Conference (WSC)*, 569–580. IEEE.

Sew, C.; Schömbs, S.; Kelly, R. M.; and Goncalves, J. 2025. User Willingness to Engage with AI-Generated Influencer Content. In *Proc. CHI Conference on Human Factors in Computing Systems*, 1–7.

Shen, J.; Mire, J.; Park, H. W.; Breazeal, C.; and Sap, M. 2024. Heart-felt narratives: Tracing empathy and narrative style in personal stories with llms. *arXiv preprint arXiv:2405.17633*.

Skulmowski, A.; and Engel-Hermann, P. 2025. The ethics of erroneous AI-generated scientific figures. *Ethics and Information Technology*, 27(2): 31.

Solano, T.; Mittal, R.; and Shoele, K. 2021. One size fits all?: A simulation framework for face-mask fit on population-based faces. *PLoS One*, 16(6): e0252143.

Solano, T.; and Shoele, K. 2022. Investigation of the role of face shape on the flow dynamics and effectiveness of face masks. *Fluids*, 7(6): 209.

Sorin, V.; Brin, D.; Barash, Y.; Konen, E.; Charney, A.; Nadkarni, G.; and Klang, E. 2024. Large language models and empathy: systematic review. *Journal of medical Internet research*, 26: e52597.

Tel, T.; and Minor, M. 2025. Utilizing the Structure of Process Models for Guided Generation of Explanatory Texts. In *Int. Conf. on Case-Based Reasoning*, 157–171. Springer.

Tomasiello, D. B.; Giannotti, M.; and Feitosa, F. F. 2020. ACCESS: An agent-based model to explore job accessibility inequalities. *Computers, Environment and Urban Systems*, 81: 101462.

Vanhée, L.; Borit, M.; Siebers, P.-O.; Cremades, R.; Frantz, C.; Gürcan, Ö.; Kalvas, F.; Kera, D. R.; Nallur, V.; Narasimhan, K.; et al. 2025. Large Language Models for Agent-Based Modelling: Current and possible uses across the modelling cycle. *arXiv preprint arXiv:2507.05723*.

Wang, C.; Sundstedt, V.; and Garro, V. 2025. Generative Artificial Intelligence for Immersive Analytics. In *20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, VISIGRAPP 2025, Porto, Feb 26-28, 2025*, volume 1, 938–946. SciTePress.

Wang, T.-K.; Solano, T.; and Shoele, K. 2022. Bridge the gap: Correlate face mask leakage and facial features with 3D morphable face models. *Journal of exposure science & environmental epidemiology*, 32(5): 735–743.

Wozniak, M. K.; and Giabbanelli, P. J. 2021. Comparing implementations of cellular automata as images: A novel approach to verification by combining image processing and machine learning. In *Proc. 2021 ACM SIGSIM conference on principles of advanced discrete simulation*, 13–25.

Xie, Y.; Wang, J.; Feng, T.; Ma, F.; and Li, Y. 2025. Ccis-diff: A generative model with stable diffusion prior for controlled colonoscopy image synthesis. In *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.

Xue, Y.; Jabi, W.; Woolley, T. E.; and Kaouri, K. 2024. Modelling indoor airborne transmission combining architectural design and people movement using the VIRIS simulator and web app. *Scientific Reports*, 14(1): 28220.

Zhang, R.; Yu, B.; Min, J.; Xin, Y.; Wei, Z.; Shi, J. N.; Huang, M.; Kong, X.; Xin, N. L.; Jiang, S.; et al. 2025. Generative AI for Film Creation: A Survey of Recent Advances. In *Proc. Computer Vision and Pattern Recognition Conference*, 6267–6279.

Zhao, Y.; Vetter, N.; and Aryan, K. 2024. Using Large Language Models for OntoClean-based Ontology Refinement. *arXiv:2403.15864*.

Zhou, H. A.; Wolfschläger, D.; Florides, C.; Werheid, J.; Behnen, H.; Woltersmann, J.-H.; Pinto, T. C.; Kemmerling, M.; Abdelrazeq, A.; and Schmitt, R. H. 2025. Generative AI in industrial machine vision: a review. *Journal of Intelligent Manufacturing*, 1–24.

Zhuang, H.; Zhang, Y.; and Liu, S. 2023. A Pilot Study of Query-Free Adversarial Attack Against Stable Diffusion. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2385–2392.