ValidAgent – An Open Repository of Validated Generative Agents for Behavioral Simulations and Multi-Agent Systems

Leon Lührs[1,2][0009-0004-5550-1387], Christian Stindt von Dohm[1,2][0009-0001-1818-7892], Anika Bittner[1][0009-0004-7090-1557], Timo Heinrich[1][0000-0001-6827-1374], and Matthias Meyer[1,2][0000-0003-2980-4670]

[1] Hamburg University of Technology, Am Schwarzenberg-Campus 4, 21073 Hamburg, Germany
[2] United Nations University Hub on Engineering to Face Climate Change, Am Schwarzenberg-Campus 3, 21073 Hamburg, Germany
leon.luehrs@tuhh.de

Abstract. Large Language Models are increasingly used to create generative agents for behavioral simulations and multi-agent systems, as they allow researchers to build entities that exhibit seemingly human-like behavior with relatively low effort. Although generative agents are frequently parameterized using empirical data, systematic empirical validation remains scarce: there is no established practice of benchmarking agent behavior against human reference data, nor a common standard for reporting validation in a way that allows comparison across studies. To our knowledge, no dedicated infrastructure exists to support the collection, documentation, and reuse of validated agents. As a result, generative agents are difficult to transfer across studies, hard to reproduce, and lack comparable validation – ultimately limiting cumulative scientific progress. In response, we present ValidAgent, an open, web-based repository that provides a common interface to browse, compare, and select from a curated collection of generative agent profiles. The platform emphasizes empirical validation, transparency, and reproducibility by requiring each agent profile to include structured documentation covering its design rationale, behavioral traits, underlying assumptions, intended scope of application, and validation results against human reference data. To demonstrate the platform's utility, we provide an initial set of empirically validated agents grounded in the die roll honesty paradigm, a widely used experimental setup for studying dishonest behavior. Building on this, we aim to foster a researcher-driven ecosystem in which generative agents can be contributed, peer-reviewed, and benchmarked. In doing so, this initiative contributes to a more transparent and cumulative simulation practices.

Keywords: Open Repository, Generative Agents, Large Language Models.

1 Introduction

With the emergence of Large Language Models (LLMs), researchers started to develop generative agents – autonomous entities designed to simulate human-like behavior with a level of flexibility that goes beyond many traditional agent architectures.[1] Unlike conventional agents, which typically rely on explicitly defined rules and fixed decision logic, generative agents leverage LLMs to enable contextual reasoning and adaptive decision-making [6, 7]. Their behavioral identity is typically defined through structured persona descriptions that guide the model's responses. Recent work emphasizes grounding this "Homo silicus" [8, p. 1] in empirical data such as interviews and experimental findings, thereby moving beyond purely synthetic modeling towards more behaviorally validated simulations [7]. As a result, generative agents hold considerable promise for simulating complex social interactions [6, 7, 9, 10], including legal reasoning [11, 12], economic decision-making [13], political negotiation [14, 15], or consumer behavior [8, 9].

Despite this potential, the development and application of generative agents remain highly fragmented. Most agents are created in isolation for specific projects. Consequently, it remains unclear to what extent agent configurations and behavioral outcomes can be replicated or transferred across studies. Recent literature recognizes the need for greater consistency in agent design and evaluation practices. Amin et al. [1, 16] emphasize the multitude of design choices – from model family and version to hyperparameters and prompting strategies – and argue for standardization. In terms of behavioral validity, Tseng et al. [17] highlight the need for expanded datasets and behavioral benchmarks grounded in real human data. However, approaches to empirically grounding generative agents vary considerably across studies and are often not documented in comparable ways. Without shared standards and infrastructure for design consistency and empirical grounding, generative agents remain difficult to compare, reproduce, and build upon, ultimately limiting cumulative scientific progress in generative agent research.

While several tools and infrastructures address individual aspects of generative agent development, to our knowledge, none provides a comprehensive solution for validated, reusable agent profiles. For instance, persona collections such as PersonaHub [18] offer large-scale, synthetically generated profiles at unprecedented scale, while DeepPersona [19] extends this approach with a more extensive attribute taxonomy, greater demographic diversity, and customizable profile generation. However, these resources focus on scalable synthetic persona generation rather than directly collected individual-level human datasets. At the same time, agent-oriented frameworks such as the GenAgents architecture [6], Concordia [20], or EDSL (Expected Parrot's Domain-Specific Language) [21] provide powerful environments for constructing, simulating, and in some cases benchmarking generative agents against human responses. Yet they are primarily

[1] There is no standardized terminology for these agents. We use "generative agents" as a working label, but note that terms such as "LLM-based agents," "generative AI," "AI agents," or "agentic AI" are also in use and may partially overlap [see, e.g., 1–5]. We conceptualize generative agent behavioral simulations as part of the broader multi-agent systems (MAS) field.

designed for agent execution and experimentation rather than for the sharing and reuse of validated agent profiles across studies. Dedicated repository infrastructures in other research domains demonstrate the value of structured sharing for cumulative and comparable research [22]. For example, platforms such as CoMSES Net [23], rbb4abm [22], AGENTBLOCKS [24, 25], OpenAI Gym [26], or Sotopia [27] demonstrate how curated repositories and benchmarking ecosystems advance research in their respective fields. Similarly, general-purpose repositories such as Zenodo [28] or OSF [29, 30] could technically host agent artifacts, yet lack domain-specific standards, structured documentation requirements, and community-driven quality control mechanisms. Platforms such as GitHub [31] or Hugging Face [32] provide versioned hosting for code, models, and datasets, but are not designed to curate generative agent profiles according to domain-specific validation standards or to treat them as structured, reviewable research objects.

In response, we propose a web-based, open-access repository dedicated to empirically validated generative agents – ValidAgent. This platform offers a common interface for researchers to share, browse, compare, and select from a curated collection of agent profiles. Each profile integrates detailed documentation encompassing design rationale, behavioral characteristics, underlying assumptions, intended application scope, and validation outcomes against human reference data. Our approach prioritizes transparency, reproducibility, and reusability to foster a research ecosystem where generative agents can be contributed, peer-reviewed, benchmarked, and cited [33]. As a first application and illustration, we provide an initial set of agents validated within the well-established die roll honesty paradigm [34].

2 A New Repository Design and Architecture – ValidAgent

The ValidAgent repository is implemented as a web-based, open-access platform that enables researchers to discover, evaluate, and export validated generative agents with minimal technical knowledge.[2] Its design is structured around two core entities: *Agent Sets* and *Agent Cards*.

Agent Sets constitute the primary unit of organization. Each set represents a curated collection of agents that were jointly validated against, e.g., a shared human reference dataset or experimental paradigm (see Section 3). The repository allows users to browse and filter Agent Sets by application context, empirical grounding, and validation approach (see Fig. 1). Within each set, more detailed information is provided, covering how the agents were constructed and for which purpose, their recommended scope of application, the components and configuration options of individual agents (e.g., demographic attributes or Big Five dimensions as inputs), validation outcomes including visualizations of representative results, as well as provenance, versioning, and licensing information.

[2] The ValidAgent repository is available at: http://www.validagent.org

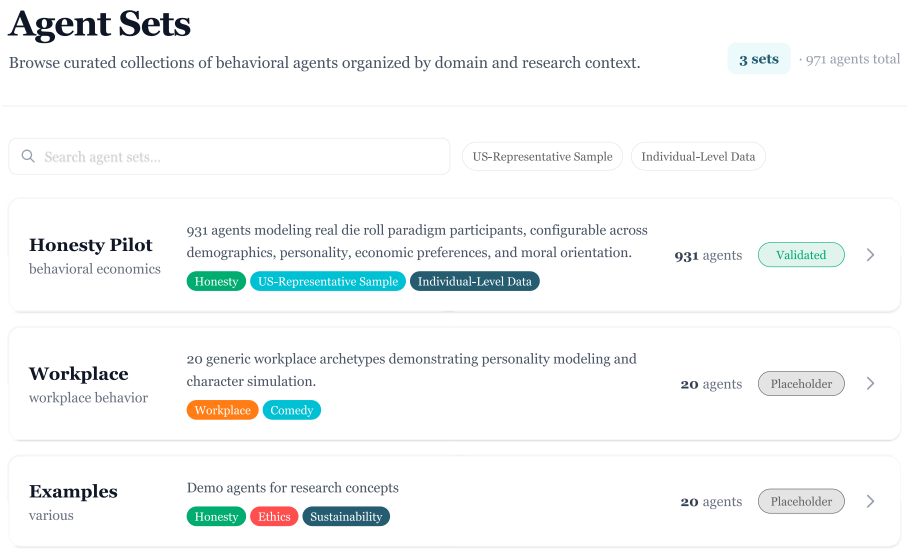

Fig. 1. Agent Set Overview.

Within each Agent Set, Agent Cards present individual agent profiles, displaying the agent inputs (e.g., a persona prompt) that serve as the direct initialization of the agent (see Fig. 2).

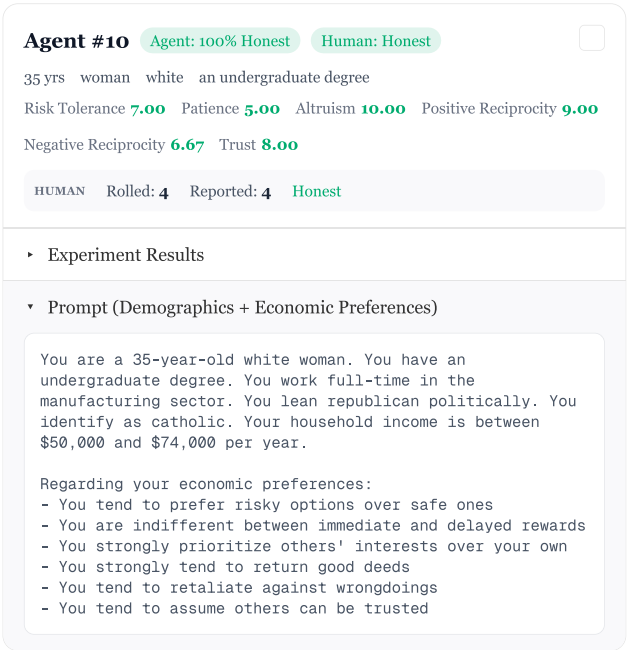

Fig. 2. Agent Card Persona Module.

Inputs are composed of modular components that can be toggled on or off, allowing researchers to tailor and export agent configurations to their needs. Where a shared configuration has been defined at the set level, the inputs are shown in accordance with the selected configuration. Where available, Agent Cards additionally display individual simulation results for the respective agent. By providing validation information at both the set and the individual agent level, the platform supports informed decisions about the appropriate scope and level of detail for a given simulation context (see Section 3).

Configured agent profiles can be exported in structured formats including JSON and YAML. To ensure interoperability, the export layer adapts configurations to the input requirements of established frameworks, including GenAgents [6], Concordia [20], EDSL [21], Sotopia [27], and OASIS [35]. Because these frameworks differ in how they initialize agents – ranging from narrative persona descriptions with memory structures to structured schemas or scenario-based specifications (see Table 1) – the platform maps its internal representation to framework-specific import formats, enabling direct integration into existing simulation environments.

Table 1. Generative Agent Simulation Frameworks.

Framework	Persona Initialization Mechanism	Import Format
GenAgents [6]	Narrative persona description + memory initialization	Background description + memory entries
Concordia [20]	Narrative persona description + memory initialization	Description string + list of fact strings
Sotopia [27]	Scenario-based character specification	Character fields embedded in scenario configuration
EDSL [21]	Structured persona schema	JSON persona input
OASIS [35]	Agent profile specification within simulated platform	Agent profile fields in environment configuration

New Agent Sets can be contributed through a structured submission interface requiring a standardized format of agent data, documentation, and accompanying metadata on validation procedures and human reference datasets. As research questions, data collection efforts, and simulations vary considerably across studies, the platform does not prescribe a single validation method. Rather, it requires that whichever empirical or experimental grounding approach is used, for example, demographic, attitudinal, or behavioral outcome benchmarks, be documented with explicit metrics. Each submission undergoes a community review process to ensure completeness of documentation, traceability of the validation process, and technical functionality. Integrated version control guarantees that each published set remains citable, while subsequent updates are handled through a transparent re-review process, ensuring long-term reproducibility [23, 24, 33].

3 Pilot Honesty Agent Set: The Die Roll Paradigm

To illustrate the repository's functionality, we introduce an empirically validated Agent Set grounded in the die roll paradigm [34] – a widely used experimental framework for studying dishonest behavior. The set includes individual agent input data, a behavioral benchmark using human reference data, and indexed simulation runs across selected input combinations [36] using GPT-5.4 mini, all browsable in the live repository.[3] It thereby serves as a reference implementation for how validated generative agents can be documented, benchmarked, and reused.

The underlying behavioral experiment is based on an online implementation of the die roll paradigm with a final representative sample of 931 US participants. In our implementation, participants rolled a six-sided virtual die and subsequently reported the outcome. Compensation depends on the reported rather than the actual outcome: higher monetary rewards are associated with reported values from one to five, while reporting a six yields no additional payment. This creates a direct financial incentive to dishonestly report a five irrespective of the actual outcome. A key advantage of our online implementation is that, unlike conventional laboratory experiment implementations using a physical die and unobservable rolls, both the actual roll outcome and the reported value can be recorded. This allows us to determine which participants reported (un)truthfully [37, 38].

Building on this setup, we additionally collected a range of persona-relevant participant characteristics[4], which were used to construct corresponding generative agent profiles. The richness of the collected data allows for flexible agent configuration by systematically varying the amount and type of information included in agent initialization [39].

At the Agent Set level Fig. 3 illustrates the honesty rate across 5 simulation runs separately for agents grounded in dishonest and honest humans. Agents, when given input information from dishonest humans under a selected input configuration, were five times more likely to choose dishonesty (20%) compared to agents based on honest humans (4%). Across both groups, the overall agent honesty rate (95%) closely approximates that observed in the human sample (94.2%).

---

[3] The pilot Agent Set was created as part of a separate empirical research project in which both human participant data and corresponding generative agent simulations were collected. The study received approval from the responsible Institutional Review Board, and informed consent was obtained from all human participants prior to participation. The repository stores the resulting agent profiles together with structured metadata, validation information, and indexed simulation outputs.

[4] To reduce the risk of including non-human responses during the human part of our study, participants were required to complete a Completely Automated Public Turing Test to Tell Computers and Humans Apart (CAPTCHA). They were also asked to select the correct image in a task where the labels assigned to the images did not match the images themselves. Additionally, we recorded their response times and disabled pasting.

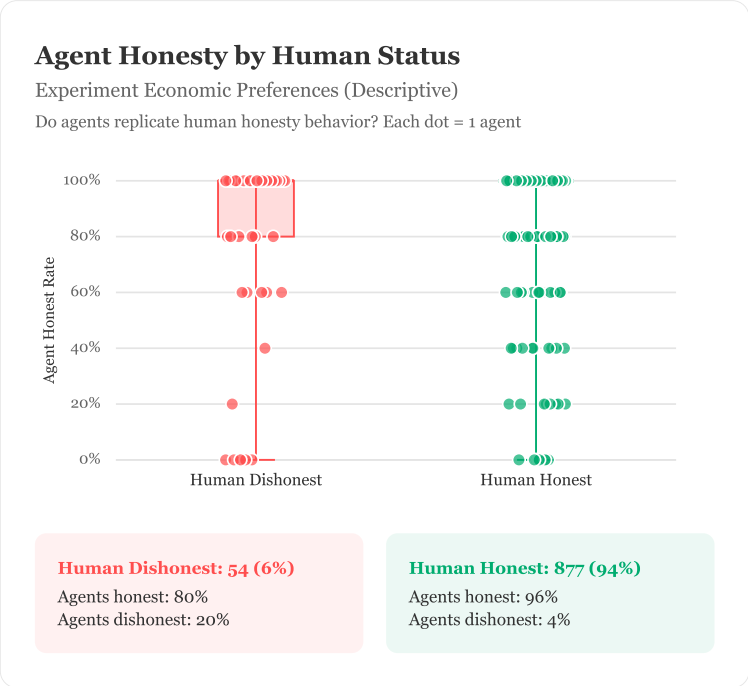

Fig. 3. Agent Set Comparison between Humans and Generative Agents.

At the Agent Card level, such a comparison can be made for individual human-agent pairs (see Fig. 4). Beyond behavioral outcomes, this set additionally enables a comparison of the reasoning underlying agent and human decisions, based on elicited explanations from both.

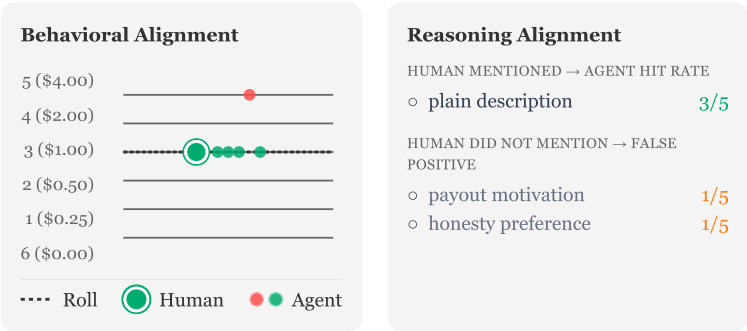

Fig. 4. Agent Card Comparison between Humans and Generative Agents.

The die roll paradigm is a particularly suitable validation domain as dishonest behavior observed in this task has been shown to correlate with real-world honesty-related behavior [40–42]. Consequently, validated agents derived from this paradigm may pro-

vide useful proxies for a broader range of honesty-related applications beyond the original experiment. These include related behavioral paradigms such as the coin-flip task [43], sender–receiver game [44, 45], or matrix task [46][5], as well as organizational and managerial contexts such as the use for fraud simulations in virtual laboratories [48, 49] or the exploration of intervention effects using empirically grounded generative agents. We encourage researchers to build upon this demo set by contributing their own validated generative agents, explore new behavioral paradigms, and empirical benchmarks.

4 Conclusion

Generative agents hold significant promise for advancing behavioral simulations and multi-agent systems [50], yet their current development landscape is fragmented, lacking shared standards for validation, documentation, and reuse. In this paper, we introduce ValidAgent, a web-based, open-access repository for empirically validated generative agents that treats agent profiles as structured, versioned, and reviewable research objects. By organizing agents into validated sets, providing modular and transparent documentation, enabling interoperability with existing agent frameworks, and establishing quality assurance mechanisms, the platform lays the groundwork for reproducible, comparable, and citable generative agent resources. In doing so, we aim to shift generative agents from isolated project-specific implementations toward a sustainable research infrastructure that supports cumulative scientific progress.

Acknowledgements. We thank the participants of the 14th International Workshop on Engineering Multi-Agent Systems (EMAS 2026) for their helpful comments. Furthermore, we are grateful to the anonymous reviewers for their constructive feedback.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

[5] For a meta-review of related honesty paradigms, see Gerlach et al. [47].

References

1. Amin, D., Salminen, J., Ahmed, F., Tervola, S.M.H., Sethi, S., Jansen, B.J.: Creating and evaluating personas using generative AI: A scoping review of 81 articles. In: Oliver, N., Shamma, D.A., Candello, H., Cesar, P., Lopes, P., Bozzon, A., Kosch, T., Liao, V., Ma, X., Artizzu, V., Draxler, F., López, G., Reinschluessel, A.V., Tong, X., and Dugas, P.O.T. (eds.) CHI '26: Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems. Article 421. Association for Computing Machinery, New York (2026). https://doi.org/https://doi.org/10.1145/3772318.379060.
2. Rio-Chanona, R.M. del, Pangallo, M., Hommes, C.: Can generative AI agents behave like humans? Evidence from laboratory market experiments (2025). https://doi.org/10.48550/arXiv.2505.07457.
3. Sapkota, R., Roumeliotis, K.I., Karkee, M.: AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges. Inf. Fusion. **126**(Part B), Article 103599 (2026). https://doi.org/10.1016/j.inffus.2025.103599.
4. Sun, G., Zhan, X., Such, J.: Building better AI agents: A provocation on the utilisation of persona in LLM-based conversational agents. In: Dubiel, M., Leiva, L.A., Trippas, J., Fischer, J., and Torre, I. (eds.) CUI '24: Proceedings of the 6th ACM Conference on Conversational User Interfaces. Association for Computing Machinery, New York (2024). https://doi.org/10.1145/3640794.3665887.
5. Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., Yin, Z., Dou, S., Weng, R., Qin, W., Zheng, Y., Qiu, X., Huang, X., Zhang, Q., Gui, T.: The rise and potential of large language model based agents: A survey. Sci. China Inf. Sci. **68**(2), 121101 (2025). https://doi.org/10.1007/s11432-024-4222-0.
6. Park, J.S., O'Brien, J., Cai, C.J., Morris, M.R., Liang, P., Bernstein, M.S., Follmer, S., Han, J., Steimle, J., Riche, N.H.: Generative agents: Interactive simulacra of human behavior. In: Follmer, S., Han, J., Steimle, J., and Riche, N.H. (eds.) UIST '23: Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology. Article 2. Association for Computing Machinery, New York (2023). https://doi.org/10.1145/3586183.3606763.
7. Park, J.S., Zou, C.Q., Kamphorst, J., Egan, N., Shaw, A., Hill, B.M., Cai, C., Morris, M.R., Liang, P., Willer, R., Bernstein, M.S.: LLM agents grounded in self-reports enable general-purpose simulation of individuals (2026). https://doi.org/10.48550/arXiv.2411.10109.
8. Horton, J.J., Filippas, A., Manning, B.S.: Large language models as simulated economic agents: What can we learn from homo silicus? https://www.nber.org/papers/w31122 (2026). https://doi.org/10.3386/w31122.
9. Li, Y., Sun, L., Zhang, Y.: MetaAgents: Large language model based agents for decision-making on teaming. In: Nichols, J. (ed.) Proceedings of the ACM on Human-Computer Interaction. Article CSCW134. Association for Computing Machinery, New York (2025). https://doi.org/10.1145/3711032.
10. Park, J.S., Popowsky, L., Cai, C., Morris, M.R., Liang, P., Bernstein, M.S.: Social simulacra: Creating populated prototypes for social computing systems. In: Agrawala, M., Wobbrock, J.O., Adar, E., and Setlur, V. (eds.) UIST '22: Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology. Article 74. Association for Computing Machinery, New York (2022). https://doi.org/10.1145/3526113.3545616.
11. Dong, Y.R., Hu, T., Collier, N.: Can LLM be a personalized judge? In: Al-Onaizan, Y., Bansal, M., and Chen, Y.-N. (eds.) Findings of the Association for Computational Linguis-

tics: EMNLP 2024. 10126–10141. Association for Computational Linguistics, Miami (2024).
12. Guha, N., Nyarko, J., Ho, D., Ré, C., Chilton, A., Narayana, A., Chohlas-Wood, A., Peters, A., Waldon, B., Rockmore, D., Zambrano, D., Talisman, D., Hoque, E., Surani, F., Fagan, F., Sarfaty, G., Dickinson, G., Porat, H., Hegland, J., Wu, J., Nudell, J., Niklaus, J., Nay, J., Choi, J., Tobia, K., Hagan, M., Ma, M., Livermore, M., Rasumov-Rahe, N., Holzenberger, N., Kolt, N., Henderson, P., Rehaag, S., Goel, S., Gao, S., Williams, S., Gandhi, S., Zur, T., Iyer, V., Li, Z.: LegalBench: A collaboratively built benchmark for measuring legal reasoning in large language models. In: Oh, A., Naumann, T., and Globerson, A. (eds.) NIPS '23: Proceedings of the 37th International Conference on Neural Information Processing Systems. 44123–44279. Curran Associates, New York (2023).
13. Li, N., Gao, C., Li, M., Li, Y., Liao, Q.: EconAgent: Large language model-empowered agents for simulating macroeconomic activities. In: Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics. 15523–15536. Association for Computational Linguistics, Bangkok (2024). https://doi.org/10.18653/v1/2024.acl-long.
14. Baker, Z.R., Azher, Z.L.: Simulating the U.S. senate: An LLM-driven agent approach to modeling legislative behavior and bipartisanship (2024). https://doi.org/10.48550/arXiv.2406.18702.
15. Hua, W., Fan, L., Li, L., Mei, K., Ji, J., Ge, Y., Hemphill, L., Zhang, Y.: War and peace (WarAgent): Large language model-based multi-agent simulation of world wars (2024). https://doi.org/10.48550/arXiv.2311.17227.
16. Amin, D., Salminen, J., Jansen, B.J., Shin, J., Kim, D.H.: Generative AI personas considered harmful? Putting forth twenty challenges of algorithmic user representation in human-computer interaction. Int. J. Hum. Comput. Stud. **205**, Article 103657 (2025). https://doi.org/10.1016/j.ijhcs.2025.103657.
17. Tseng, Y.-M., Huang, Y.-C., Hsiao, T.-Y., Chen, W.-L., Huang, C.-W., Meng, Y., Chen, Y.-N.: Two tales of persona in LLMs: A survey of role-playing and personalization. In: Al-Onaizan, Y., Bansal, M., and Chen, Y.-N. (eds.) Findings of the Association for Computational Linguistics: EMNLP 2024. 16612–16631. Association for Computational Linguistics, Miami (2024). https://doi.org/10.18653/v1/2024.findings-emnlp.969.
18. Ge, T., Chan, X., Wang, X., Yu, D., Mi, H., Yu, D.: Scaling synthetic data creation with 1,000,000,000 personas (2025). https://doi.org/10.48550/arXiv.2406.20094.
19. Wang, Z., Zhou, Y., Luo, Z., Ye, L., Wood, A., Yao, M., Mansour, S., Pan, L.: DeepPersona: A generative engine for scaling deep synthetic personas (2025). https://doi.org/10.48550/arXiv.2511.07338.
20. Vezhnevets, A.S., Agapiou, J.P., Aharon, A., Ziv, R., Matyas, J., Duéñez-Guzmán, E.A., Cunningham, W.A., Osindero, S., Karmon, D., Leibo, J.Z.: Generative agent-based modeling with actions grounded in physical, social, or digital space using Concordia (2023). https://doi.org/10.48550/arXiv.2312.03664.
21. Expected Parrot: Expected Parrot: Tools for AI-powered research, https://docs.expectedparrot.com/en/latest, last accessed 2025/12/15.
22. Berger, U., Bell, A., Barton, C.M., Chappin, E., Dreßler, G., Filatova, T., Fronville, T., Lee, A., van Loon, E., Lorscheid, I., Meyer, M., Müller, B., Piou, C., Radchuk, V., Roxburgh, N., Schüler, L., Troost, C., Wijermans, N., Williams, T.G., Wimmler, M.-C., Grimm, V.: Towards reusable building blocks for agent-based modelling and theory development. Environ. Model. Softw. **175**, Article 106003 (2024). https://doi.org/10.1016/j.envsoft.2024.106003.

23. Rollins, N.D., Barton, C.M., Bergin, S., Janssen, M.A., Lee, A.: A computational model library for publishing model documentation and code. Environ. Model. Softw. **61**, 59–64 (2014). https://doi.org/10.1016/j.envsoft.2014.06.022.
24. Filatova, T., Verbeek, L., Warnier, M., Ghorbani, A., Nikolic, I., Grimm, V., Berger, U., Barton, M., Bell, A., Lee, A., Magliocca, N.R., Wagenblast, T.: AGENTBLOCKS: A community platform for sharing, comparing, and improving reusable building blocks for (agent-based) models. J. Artif. Soc. Soc. Simul. **28**(4), Article 11 (2025). https://doi.org/10.18564/jasss.5831.
25. Kunkel, J., Grimm, V., Stindt von Dohm, C., Grothmann, L., Schwarz, L., Hotz, R., Huber, R., Zeng, Y., Jurak, K., Berger, U., Polhill, J.G., Verbeek, L., Filatova, T., Meyer, M., Müller, B., Groeneveld, J.: Fourteen reusable building blocks (RBBs) for representing decision-making in agent- based models: A compilation on AGENTBLOCKS demonstrating best practices and ways forward (2026).
26. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: OpenAI Gym (2016). https://doi.org/10.48550/arXiv.1606.01540.
27. Zhou, X., Zhu, H., Mathur, L., Zhang, R., Yu, H., Qi, Z., Morency, L.-P., Bisk, Y., Fried, D., Neubig, G., Sap, M.: SOTOPIA: Interactive evaluation for social intelligence in language agents (2024). https://doi.org/10.48550/arXiv.2310.11667.
28. Zenodo: https://zenodo.org/, last accessed 2026/02/14.
29. Soderberg, C.K.: Using OSF to share data: A step-by-step guide. Adv. Methods Pract. Psychol. Sci. **1**(1), 115–120 (2018). https://doi.org/10.1177/2515245918757689.
30. Sullivan, I., DeHaven, A., Mellor, D.: Open and reproducible research on open science framework. Curr. Protoc. Essent. Lab. Tech. **18**(1), Article e32 (2019). https://doi.org/10.1002/cpet.32.
31. GitHub: https://github.com/, last accessed 2026/02/20.
32. Hugging Face: https://huggingface.co/, last accessed 2026/02/20.
33. Wilkinson, M.D., Dumontier, M., Aalbersberg, Ij.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L.B., Bourne, P.E., Bouwman, J., Brookes, A.J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C.T., Finkers, R., Gonzalez-Beltran, A., Gray, A.J.G., Groth, P., Goble, C., Grethe, J.S., Heringa, J., 't Hoen, P.A.C., Hooft, R., Kuhn, T., Kok, R., Kok, J., Lusher, S.J., Martone, M.E., Mons, A., Packer, A.L., Persson, B., Rocca-Serra, P., Roos, M., van Schaik, R., Sansone, S.-A., Schultes, E., Sengstag, T., Slater, T., Strawn, G., Swertz, M.A., Thompson, M., van der Lei, J., van Mulligen, E., Velterop, J., Waagmeester, A., Wittenburg, P., Wolstencroft, K., Zhao, J., Mons, B.: The FAIR guiding principles for scientific data management and stewardship. Sci Data. **3**(1), Article 160018 (2016). https://doi.org/10.1038/sdata.2016.18.
34. Fischbacher, U., Föllmi-Heusi, F.: Lies in disguise: An experimental study on cheating. J. Eur. Econ. Assoc. **11**(3), 525–547 (2013). https://doi.org/10.1111/jeea.12014.
35. Yang, Z., Zhang, Z., Zheng, Z., Jiang, Y., Gan, Z., Wang, Z., Ling, Z., Chen, J., Ma, M., Dong, B., Gupta, P., Hu, S., Yin, Z., Li, G., Jia, X., Wang, L., Ghanem, B., Lu, H., Lu, C., Ouyang, W., Qiao, Y., Torr, P., Shao, J.: OASIS: Open agent social interaction simulations with one million agents (2025). https://doi.org/10.48550/arXiv.2411.11581.
36. Klingert, F.M.A., Meyer, M.: Effectively combining experimental economics and multi-agent simulation: suggestions for a procedural integration with an example from prediction markets research. Comput. Math. Organ. Theory. **18**(1), 63–90 (2012). https://doi.org/10.1007/s10588-011-9098-2.
37. Leib, M.: People imitate others' dishonesty but do not intentionally search information about it. J. Behav. Decis. Mak. **36**(2), Article e2296 (2023). https://doi.org/10.1002/bdm.2296.

38. Pascual-Ezama, D., Prelec, D., Muñoz, A., Gil-Gómez de Liaño, B.: Cheaters, liars, or both? A new classification of dishonesty profiles. Psychol Sci. **31**(9), 1097–1106 (2020). https://doi.org/10.1177/0956797620929634.
39. Lorscheid, I., Meyer, M.: Divide and conquer: Configuring submodels for valid and efficient analyses of complex simulation models. Ecol. Model. **326**, 152–161 (2016). https://doi.org/10.1016/j.ecolmodel.2015.11.013.
40. Dai, Z., Galeotti, F., Villeval, M.C.: Cheating in the lab predicts fraud in the field: An experiment in public transportation. Manage. Sci. **64**(3), 1081–1100 (2018). https://doi.org/10.1287/mnsc.2016.2616.
41. Gächter, S., Schulz, J.F.: Intrinsic honesty and the prevalence of rule violations across societies. Nature. **531**(7595), 496–499 (2016). https://doi.org/10.1038/nature17160.
42. Olsen, A.L., Hjorth, F., Harmon, N., Barfort, S.: Behavioral dishonesty in the public sector. J Public Adm Res Theory. **29**(4), 572–590 (2019). https://doi.org/10.1093/jopart/muy058.
43. Bucciol, A., Piovesan, M.: Luck or cheating? A field experiment on honesty with children. J. Econ. Psychol. **32**(1), 73–78 (2011). https://doi.org/10.1016/j.joep.2010.12.001.
44. Gneezy, U.: Deception: The role of consequences. Am. Econ. Rev. **95**(1), 384–394 (2005). https://doi.org/10.1257/0002828053828662.
45. Sánchez-Pagés, S., Vorsatz, M.: An experimental study of truth-telling in a sender–receiver game. Games Econ. Behav. **61**(1), 86–112 (2007). https://doi.org/10.1016/j.geb.2006.10.014.
46. Mazar, N., Amir, O., Ariely, D.: The dishonesty of honest people: A theory of self-concept maintenance. J. Mark. Res. **45**(6), 633–644 (2008). https://doi.org/10.1509/jmkr.45.6.633.
47. Gerlach, P., Teodorescu, K., Hertwig, R.: The truth about lies: A meta-analysis on dishonest behavior. Psychol. Bull. **145**(1), 1–44 (2019). https://doi.org/10.1037/bul0000174.
48. Eckert, A., Meyer, M., Stindt, C.: Combining vignette surveys with agent-based modeling: Insights on fraud dynamics with empirically calibrated norm sensitivities. In: Czupryna, M., Kamiński, B., and Verhagen, H. (eds.) Proceedings of the 19th Social Simulation Conference, Cracow, Poland, 16–20 September 2024. 83–98. Springer International Publishing, Cham, Switzerland (2025). https://doi.org/10.1007/978-3-031-91782-0_6.
49. Meyer, M., Stindt von Dohm, C.: Studying management control in the virtual lab: Agent-based models, cultural control, and social norms. In: Bellora-Bienengräber, L. and Weißenberger, B.E. (eds.) Handbook on Management Control in a Transforming Society. Edward Elgar Publishing (in press).
50. Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N.V., Wiest, O., Zhang, X.: Large language model based multi-agents: A survey of progress and challenges. In: Larson, K. (ed.) Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24). 8048–8057. International Joint Conferences on Artificial Intelligence, Jeju (2024).