Playing Domains: Codes, Cities, and Cultures in the Viral World of Machine Learning

Sam Hind

What happens when cities become datasets for AI competitions? Sam Hind shows how machine learning's scoreboards distance practitioners from the real-world impacts of their work.

[Ed. note: This article is part of a dossier on Playable Cities.]

The history of artificial intelligence (AI) is also a history of play. In 1988, acclaimed computer scientist Raj Reddy proposed six "grand challenges" intended to jolt the AI community into life.[1]Raj Reddy, "Foundations and Grand Challenges of Artificial Intelligence," *AI Magazine* 9, no. 4 (1988): 9–21, https://doi.org/10.1609/aimag.v9i4.950. The most famous was the goal of building a "world champion chess machine." After losing the first match in February 1996, IBM's Deep Blue computer beat reigning world chess champion Garry Kasparov the following May. More recently, AI start-up DeepMind – acquired by Google in 2014 – built AlphaGo, a computer program capable of playing Go, a game infinitely more complex than chess. After winning sixty straight online games against professionals an updated version of AlphaGo (referred to as "Master") beat the number one ranked Go player in the world, China's Ke Jie, in a three-game match – twenty years since Deep Blue's victory. Realizing there was nothing else left to win, DeepMind retired all versions of AlphaGo to "throw their considerable energy into the next set of grand challenges…such as finding new cures for diseases."[2]Sam Byford, "AlphaGo Retires from Competitive Go after Defeating World Number One 3-0," *The Verge*, May 27, 2017, https://www.theverge.com/2017/5/27/15704088/alphago-ke-jie-game-3-result-retires-future.

Yet AI's connection to play goes much deeper. For twenty years competitions, often referred to as *challenges*, have been organized to drive the development, and more recent commercialization, of AI and machine learning (ML).[3]Sam Hind, Fernando van der Vlist, and Max Kanderske, "Challenges as Catalysts: How Waymo's Open Dataset Challenges Shape AI Development," *AI & Society* 40, no. 0 (2024): 1667–1683, https://doi.org/10.1007/s00146-024-01927-x. In this article I draw on recent scholarship on "AI as a sport" to argue that as everyday urban environments – people, streets, situations – are compressed into training datasets, they have begun to function as game worlds for practitioners, attracted by the viral nature of machine learning.[4]Will Orr and Edward B. Kang, "AI as a Sport: On the Competitive Epistemologies of Benchmarking," in *FAccT '24: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1875–84, https://doi.org/10.1145/3630106.3659012. This practice of "playing domains" – where practitioners treat AI as an arena for gladiatorial-style battles – is becoming a central feature of contemporary machine learning culture driven by this virality.

AI as a sport of conflicting codes

To evaluate the performance of ML models, the AI community engage in the practice of benchmarking: a process that allows developers to assess the performance of their own models by comparing them against the performance of other models. For this comparative evaluation to take place, certain standards are established. Firstly, models need to be tested on the same *tasks*. In computer vision (a subfield of AI) these tasks might include semantic segmentation ("a computer vision task that assigns a class label to pixels"),[5]"What Is Semantic Segmentation?," IBM, accessed May 1, 2025, https://www.ibm.com/think/topics/semantic-segmentation. or object detection ("a technique that uses neural networks to localize and classify objects in an image").[6]Jacob Murel and Eda Kavlakoğlu, *What Is Object Detection?*, IBM, 2025, https://www.ibm.com/think/topics/object-detection (accessed May 1, 2025).

Secondly, models need to be tested with the same *dataset*. If used by enough models these datasets morph into benchmark datasets: “the standard that both brings together particular subcommunities of ML researchers, as well as further enables their ‘progression’ through the competitive and iterative comparison of ML models.”[7]Orr and Kang, “AI as a Sport,” 1876. For Isak Engdahl, “the process of crafting benchmark datasets involves adherence to agreed-upon standards,” before they can be used as standards in and of themselves.[8]Isak Engdahl, “Agreements ‘in the Wild’: Standards and Alignment in Machine Learning Benchmark Dataset Construction,” *Big Data & Society* 11, no. 2 (2024): 1–13, https://doi.org/10.1177/20539517241242457. Much of the work involved in creating benchmark datasets requires what Engdahl refers to as “alignment work,”[9]Engdahl, “Agreements ‘in the Wild’,” 4. sets of agreed-upon conventions, rules, and guidelines that help to standardize data collection practices such that a “benchmark standard” can be achieved.[10]Engdahl, “Agreements ‘in the Wild’,” 9.

In an urban context, the establishment of the KITTI Vision Benchmark Suite was a critical moment in the development of autonomous vehicles,[11]Sam Hind, *Driving Decisions: How Autonomous Vehicles Make Sense of the World* (Palgrave Macmillan, 2024). driven by an identified “lack of demanding benchmarks that [could] mimic” the messiness of real driving worlds.[12]Andreas Geiger, Philip Lenz, and Raquel Urtasun, “Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite,” in *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3354–61, 2012, https://doi.org/10.1109/CVPR.2012.6248074, 3354. Without the real-world images – originally, 194 training scenes captured in Karlsruhe, Germany using a stereo camera – autonomous vehicles were destined to never leave the robotics lab. In technical parlance it needed more than a statistical “ground truth.”[13]Florian Jaton, *The Constitution of Algorithms: Ground-Truthing, Programming, Formulating* (MIT Press, 2021). Put plainly, it needed a real city.

To squeeze the most out of this compressed city, metrics such as Mean Average Precision (mAP) are used “to analyze the performance of object detection and segmentation systems” and leaderboards enable the ranking of models tested on leading benchmarks like the KITTI Vision Benchmark Suite.[14]Deval Shah, “Mean Average Precision (mAP) Explained: Everything You Need to Know,” *V7 Labs*, March 7, 2022, https://www.v7labs.com/blog/mean-average-precision. Together models, datasets, metrics, and leaderboards comprise the key ingredients in the “competitive epistemologies of benchmarking,” written deep into the culture, psyche, and practices of the AI community.[15]Orr and Kang, “AI as a Sport,” 1875.

Yet the game of building the ‘“best” model is fraught with unknowns: victories contested, rules bent, leaderboards gamed. Even benchmark datasets are flawed, precisely because cities are full of quirks and abnormalities. What works in Karlsruhe doesn’t in San Francisco.[16]Andrew Hawkins, “Cruise says a hit-and-run ‘launched’ pedestrian in front of one of its robotaxis,” *The Verge*, October 3, 2023, https://www.theverge.com/2023/10/3/23901233/cruise-crash-hit-run-pedestrian-injury-sf-robotaxi. Training models on established benchmarks may enable teams to top leaderboards, but leaderboards are poor reflections of real-world utility.[17]Kawin Ethayarajh and Dan Jurafsky, “Utility is in the Eye of the User: A Critique of NLP Leaderboards,” *arXiv* (2021): 1–8, https://doi.org/10.48550/arXiv.2009.13888. Instead, AI may be better understood as a sport of conflicting codes played, followed, or ignored, at any one given time. Even though the competitive epistemologies of AI have driven various kinds of standardization work designed to better “align” developer interests and goals within the AI community, sporting chaos still reigns.

Machine learning is a city

One way the AI community has found to straighten out these conflicts is through the organization of challenges.[18]Dieuwertje Luitse, Tobias Blanke, and Thomas Poell, “AI Competitions as Infrastructures of Power in Medical Imaging,” *Information, Communication & Society* 0, no. 0 (2024): 1–23, https://doi.org/10.1080/1369118X.2024.2334393, Hind et al., “Challenges as Catalysts”. Increasingly hosted on dedicated online platforms such as Kaggle or Grand Challenge, these competitions actualize the competitive epistemologies of benchmarking, which otherwise float freely throughout the AI community.[19]“Competitions,” Kaggle, accessed May 1, 2025, https://www.kaggle.com/competitions. In a world overflowing with models, challenges seek to focus and commodify "scholarly attention" in the discipline.[20]Samuel Goree et al., "Attention Is All They Need: Exploring the Media Archaeology of the Computer Vision Research Paper," *Proceedings of the ACM on Human-Computer Interaction* 8, no. CSCW2 (2024): 1–25, https://doi.org/10.1145/3686955. Participation in such challenges is highly valued within the AI community, and success can bring a myriad of opportunities, from direct cash prizes to future employment at the tech firms and start-ups responsible for organizing challenges.[21]Hind et al., "Challenges as Catalysts".

Whilst challenges have historically been organized by computer scientists for the benefit of the scientific community,[22]Mark Everingham et al., "The PASCAL Visual Object Classes Challenge: A Retrospective," *International Journal of Computer Vision* 111 (2015): 98–136, https://doi.org/10.1007/s11263-014-0733-5. they have also been organized by companies: US streaming behemoth Netflix offered $1 million to anyone who could build a better algorithmic recommendation system than their own.[23]Blake Hallinan and Ted Striphas, "Recommended for You: The Netflix Prize and the Production of Algorithmic Culture," *New Media & Society* 18, no. 1 (2016): 117–37, https://doi.org/10.1177/1461444814538646. One of the perennial difficulties for AI firms is knowing how to find "talent" – freshly-qualified computer science graduates. Due to the huge costs in training models and running computational hardware, "companies with the cash to spend on the best talent are the ones who tend to be winning the race" for it.[24]Srnicek, "Data, Compute, Labor," 252. Competitions make that process a whole lot easier – and cheaper.

Challenges are facilitated through the widespread availability of open-source software (OSS) for building and testing ML models. Rather than alternatives to proprietary software, open systems such as Google's Tensorflow or Meta's PyTorch enable these firms to shape and underpin the development of AI products.[25]David Gray Widder, Meredith Whittaker, and Sarah Myers West, "Why 'Open' AI Systems Are Actually Closed, and Why This Matters," *Nature* 635 (2024): 827–33, https://doi.org/10.1038/s41586-024-08141-1. As Nick Srnicek has written, "the seemingly noncapitalist practice of releasing their AI software for free in fact obscures a significant capitalist battle between…major companies."[26]Nick Srnicek, "Data, Compute, Labor," in *Digital Work in the Planetary Market*, ed. Mark Graham and Fabian Ferrari (MIT Press, 2022), 252. With new users locked-in, software functions as "feeder networks for the emerging generations of talent" in the AI industry.[27]Srnicek, "Data, Compute, Labor," 253.

But what happens when practitioners decide to take these firms up on their generous offer? Through their availability, ease-of-use, and interoperability OSS function akin to "sandbox computer games… where the player can experiment and explore in an open-ended fashion."[28]Peter Nelson, "Claustrophia, Repetition and Redundancy: The Economy and Aesthetics of User-Generated Content in Sandbox Computer Games," *Game Studies: The International Journal of Computer Game Research* 23, no. 2 (2023), https://gamestudies.org/2302/articles/nelson. Aided by online training courses, practitioners are able to acquire necessary skills in ML techniques without additional need for formal training.[29]Inga Luchs, Clemens Apprich, and Marcel Broersma, "Learning Machine Learning: On the Political Economy of Big Tech's Online AI Courses," *Big Data & Society* 10, no. 1 (2023): 1–12, https://doi.org/10.1177/20539517231153806. Although challenges typically involve strict rules of participation, open ML systems offers a space for non-linear experimentation: a kind of Minecraft for the AI community.

This combination of challenges and sandbox environments offers a "vibrant ecosystem" for experts and enthusiasts alike.[30]"OpenML: A Worldwide Machine Learning Lab," OpenML, accessed September 1, 2025, https://www.openml.org/. Whilst play "is a fundamental way of obtaining knowledge," and an "important part of socio-technical change," as Glas and Lammes have contended,[31]René Glas and Sybille Lammes, "Ludo-Epistemology: Playing with the Rules in Citizen Science Games," in *The Playful Citizen: Civic Engagement in a Mediatized Culture*, ed. René Glas et al. (Amsterdam University Press, 2019), 220. this amalgam of tools and software turns Shannon Mattern's maxim on its head: the city may not be a computer, but machine learning culture *is* a city, styled as a *demos*, home to those who want to "build a more inclusive and frictionless ecosystem of data, tools and clear results" or participate in a "mission to democratize good machine learning, one commit at a time."[32]Shannon Mattern, *A City is Not a*

Computer: Other Urban Intelligences (Princeton University Press, 2021), "Contribute," OpenML, accessed September 1, 2025, https://www.openml.org/contribute, "Hugging Face," Hugging Face, accessed September 1, 2025, https://huggingface.co/huggingface.

But the firms who build open-source resources also benefit from "communities of labor that provide inputs back into the software – all for free,"[33]Srnicek, "Data, Compute, Labor," 252. not dissimilar to how the computer game "modding" community was "used as a recruiting pool for the games industry" during the 2000s.[34]Julian Kücklich, "Precarious Playbour: Modders and the Digital Games Industry," *The Fibreculture Journal* 5 (2005), https://five.fibreculturejournal.org/fcj-025-precarious-playbour-modders-and-the-digital-games-industry/. In such a case – much like in machine learning today – "the importance of this 'free' source of innovation can hardly be overstated."[35]Kücklich, "Precarious Playbour".

Viral culture

The consequence is that the AI community engage in a practice of *playing domains* stimulated by an assortment of models, benchmarks, competitions and leaderboards. Attracted by the bright lights of machine learning's big city, practitioners become hooked by viral tasks and competitions alike.[36]Ethayarajh and Jurafsky, "Utility is in the Eye of the User: A Critique of NLP Leaderboards". ML platforms like Hugging Face, LM Arena, and OpenRouter each host public leaderboards that further boost engagement, interaction, and creative experimentation with different kinds of models like the Large Language Models (LLMs) driving generative AI's boom.[37]"Open LLM Leaderboard," Hugging Face, accessed September 1, 2025, https://huggingface.co/spaces/open-llm-leaderboard, "Leaderboard Overview," LMArena, accessed September 1, 2025, https://lmarena.ai/leaderboard, "Rankings," OpenRouter, accessed September 1, 2025, https://openrouter.ai/rankings.

Internal drivers still remain: like the desire to propose "elegant solutions" to unresolved technical problems.[38]Erica Klarreich, "A Decades-Old Computer Science Puzzle Was Solved in Two Pages," August 4, 2019, https://www.wired.com/story/a-decades-old-computer-science-puzzle-was-solved-in-two-pages/. But although the competitive epistemologies of benchmarking are governed by the "golden ethic" of comparative technical performance,[39]Sam Hind and Tatjana Seitz, "Cynical Technical Practice: From AI to APIs," *Convergence: The International Journal of Research into New Media Technologies* 30, no. 1 (2022): 29–48, https://doi.org/10.1177/13548565221133248. in the increasingly viral world of machine learning interest is piqued through the attractive properties of any one model, benchmark, task or competition rather than the inherent properties of a domain. Whilst each applied domain – be it autonomous driving, medical imaging, or facial recognition – might still be treated as a priori equal, interest in them is not.

This can be considered the practical effect of what Ribes et al. refer to as the "logic of domains," a "key organizing principle for contemporary computing projects and in broader science policy."[40]David Ribes et al., "The Logic of Domains," *Social Studies of Science* 49, no. 3 (2019): 281, https://doi.org/10.1177/0306312719849709. In such instances domains function as a "real world" to AI practitioners,[41]Ribes et al., "The Logic of Domains," 285. a place simply where (recently-acquired) foundational knowledge can be tested.[42]Adrian Mackenzie, *Machine Learners: Archaeology of a Data Practice* (MIT Press, 2017). Boosted by their viral status, challenges function as "structuring devices" that facilitate this translation of knowledge, from lab to domain.[43]Hind et al., "Challenges as Catalysts," 5.

This is the real consequence of AI as a sport: that the worlds in which such AI innovations are meant to improve in some technical way, come second to the allure of ML competitions with $1 million prize funds.[44]"The State of Machine Learning Competitions: 2024 Edition," ML Contests, accessed September 1, 2025, https://mlcontests.com/state-of-machine-learning-competitions-2024/. Moreover, that this indifference – rather than an overtly moral position or principle – is driven by the viral nature of contemporary machine learning culture. Here, ML practitioners are thrice removed from the impact of their work in the urban environments in which they are designed to be implemented: firstly, from the phenomena encoded in benchmark datasets (people, streets, urban life), secondly, from the statistical relations established in ML models, and thirdly, from situations and scenarios backgrounded in competitions. Challenge entrants do not need to understand, or care about, the people, places, or phenomena behind their AI work, only how their models perform when compared to others. Play, and the playing of domains, is at the heart of the quest for fame and glory, "an indispensable feature of machine learning culture."[45]Orr and Kang, "AI as a Sport," 1876.

Conclusion

If AI is a sport, who are the spectators?[46]Thanks to Markus Stauff for posing this question during an invited talk in the ASCA Cities seminar series on 'Playable Cities' at the University of Amsterdam (NL) in March 2025. The answer depends. Challenges invite everyone to join the game itself – tools are free, datasets available, playing field level. No one need stand on the sidelines – unless you happen to be a resident of a country subject to US sanctions.[47]To be eligible to compete in the Waymo Open Dataset Challenges in 2023, entrants, for instance, could not be a 'resident of Brazil, Quebec, Italy, or any country, state, province or territory subject to comprehensive OFAC [US Office of Foreign Assets Control] sanctions, including Cuba, Iran, North Korea, Syria, or the regions of Crimea, Donetsk or Luhansk of Ukraine,' "WOD Challenges Official Rules," Waymo, accessed May 1, 2025, https://waymo.com/open/terms/. As I have also argued here, the urban figures variously in the viral world of machine learning. Firstly, cities become compressed into training datasets used in benchmarking ML models. Secondly, the explosion of free ML software and tools constitutes machine learning as a city itself, vibrating with possibilities. Then thirdly, domains in which AI might nominally become useful within the city are impacted by the viral attractiveness of ML competitions promising fame, glory, and a cash prize.

Beneath the supposed egalitarianism of AI as a sport, not everything is equal. Computation is the great divider determining the favourites and likely winners. Experimenting in sandbox environments is only risk-free to a point: Google Cloud Credits used to train models aren't *gratis* forever.[48]"Free Cloud Features and Trial Offer," Google Cloud, accessed May 1, 2025, https://cloud.google.com/free/docs/free-cloud-features. Challenges cultivate a competitive world driven by investment capability alone, only constrained by light-touch regulation and occasional antitrust rulings.[49]Jess Weatherbed and Lauren Felner, "Google Search Charged with Breaking EU Antitrust Rules," *The Verge*, March 19, 2025, https://www.theverge.com/news/618168/google-search-eu-dma-charge-violation. The presumed winners are those deemed to have built a leaderboard-topping model. The real winners are challenge organizers and software providers – the ones who make the rules, provide the tools, and profit from the fruits of participants' playbour.[50]Kücklich, "Precarious Playbour".

This provokes a follow-up question: if AI is a sport, who loses? One answer is any "expert and enthusiast alike" incapable of building a high-performing, leaderboard-topping, competition-winning model. As the world of machine learning is built on virality and attention, those that cannot generate it by winning competitions – or spend enough time telling people they did – lose out.[51]Goree et al., "Attention is All They Need". Another answer is perhaps truer still: the AI community itself, deprived of an even more "vibrant ecosystem" of rule designers, tool-makers, method-inventors, and game-players as some predicted.[52]Everingham et al., "The PASCAL Visual Object Classes Challenge". The ultimate losers, arguably those subject to AI's viral culture, are those out there in the real world, spectators to AI community's thirst for competition.

Notes

Notes

↑1 Raj Reddy, "Foundations and Grand Challenges of Artificial Intelligence," *AI Magazine* 9, no. 4 (1988): 9–21, https://doi.org/10.1609/aimag.v9i4.950.

↑2 Sam Byford, "AlphaGo Retires from Competitive Go after Defeating World Number One 3-0," *The Verge*, May 27, 2017, https://www.theverge.com/2017/5/27/15704088/alphago-ke-jie-game-3-result-retires-future.

↑3 Sam Hind, Fernando van der Vlist, and Max Kanderske, "Challenges as Catalysts: How Waymo's Open Dataset Challenges Shape AI Development," *AI & Society* 40, no. 0 (2024): 1667–1683, https://doi.org/10.1007/s00146-024-01927-x.

↑4 Will Orr and Edward B. Kang, "AI as a Sport: On the Competitive Epistemologies of Benchmarking," in *FAccT '24: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1875–84, https://doi.org/10.1145/3630106.3659012.

↑5 "What Is Semantic Segmentation?," IBM, accessed May 1, 2025, https://www.ibm.com/think/topics/semantic-segmentation.

↑6 Jacob Murel and Eda Kavlakoğlu, *What Is Object Detection?*, IBM, 2025, https://www.ibm.com/think/topics/object-detection (accessed May 1, 2025).

↑7, ↑45 Orr and Kang, "AI as a Sport," 1876.

↑8 Isak Engdahl, "Agreements 'in the Wild': Standards and Alignment in Machine Learning Benchmark Dataset Construction," *Big Data & Society* 11, no. 2 (2024): 1–13, https://doi.org/10.1177/20539517241242457.

↑9 Engdahl, "Agreements 'in the Wild'," 4.

↑10 Engdahl, "Agreements 'in the Wild'," 9.

↑11 Sam Hind, *Driving Decisions: How Autonomous Vehicles Make Sense of the World* (Palgrave Macmillan, 2024).

↑12 Andreas Geiger, Philip Lenz, and Raquel Urtasun, "Are We Ready for Autonomous Driving? The KITTI Vision Benchmark Suite," in *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3354–61, 2012, https://doi.org/10.1109/CVPR.2012.6248074, 3354.

↑13 Florian Jaton, *The Constitution of Algorithms: Ground-Truthing, Programming, Formulating* (MIT Press, 2021).

↑14 Deval Shah, "Mean Average Precision (mAP) Explained: Everything You Need to Know," *V7 Labs*, March 7, 2022, https://www.v7labs.com/blog/mean-average-precision.

↑15 Orr and Kang, "AI as a Sport," 1875.

↑16 Andrew Hawkins, "Cruise says a hit-and-run 'launched' pedestrian in front of one of its robotaxis," *The Verge*, October 3, 2023, https://www.theverge.com/2023/10/3/23901233/cruise-crash-hit-run-pedestrian-injury-sf-robotaxi.

↑17 Kawin Ethayarajh and Dan Jurafsky, "Utility is in the Eye of the User: A Critique of NLP Leaderboards," *arXiv* (2021): 1–8, https://doi.org/10.48550/arXiv.2009.13888.

↑18 Dieuwertje Luitse, Tobias Blanke, and Thomas Poell, "AI Competitions as Infrastructures of Power in Medical Imaging," *Information, Communication & Society* 0, no. 0 (2024): 1–23, https://doi.org/10.1080/1369118X.2024.2334393, Hind et al., "Challenges as Catalysts".

↑19 "Competitions," Kaggle, accessed May 1, 2025, https://www.kaggle.com/competitions.

↑20 Samuel Goree et al., "Attention Is All They Need: Exploring the Media Archaeology of the Computer Vision Research Paper," *Proceedings of the ACM on Human-Computer Interaction* 8, no. CSCW2 (2024): 1–25, https://doi.org/10.1145/3686955.

↑21 Hind et al., "Challenges as Catalysts".

↑22 Mark Everingham et al., "The PASCAL Visual Object Classes Challenge: A Retrospective," *International Journal of Computer Vision* 111 (2015): 98–136, https://doi.org/10.1007/s11263-014-0733-5.

↑23 Blake Hallinan and Ted Striphas, "Recommended for You: The Netflix Prize and the Production of Algorithmic Culture," *New Media & Society* 18, no. 1 (2016): 117–37, https://doi.org/10.1177/1461444814538646.

↑24, ↑33 Srnicek, "Data, Compute, Labor," 252.

↑25 David Gray Widder, Meredith Whittaker, and Sarah Myers West, "Why 'Open' AI Systems Are Actually Closed, and Why This Matters," *Nature* 635 (2024): 827–33, https://doi.org/10.1038/s41586-024-08141-1.

↑26 Nick Srnicek, "Data, Compute, Labor," in *Digital Work in the Planetary Market*, ed. Mark Graham and Fabian Ferrari (MIT Press, 2022), 252.

↑27 Srnicek, "Data, Compute, Labor," 253.

↑28 Peter Nelson, "Claustrophia, Repetition and Redundancy: The Economy and Aesthetics of User-Generated Content in Sandbox Computer Games," *Game Studies: The International Journal of Computer Game Research* 23, no. 2 (2023), https://gamestudies.org/2302/articles/nelson.

↑29 Inga Luchs, Clemens Apprich, and Marcel Broersma, "Learning Machine Learning: On the Political Economy of Big Tech's Online AI Courses," *Big Data & Society* 10, no. 1 (2023): 1–12, https://doi.org/10.1177/20539517231153806.

↑30 "OpenML: A Worldwide Machine Learning Lab," OpenML, accessed September 1, 2025, https://www.openml.org/.

↑31 René Glas and Sybille Lammes, "Ludo-Epistemology: Playing with the Rules in Citizen Science Games," in *The Playful Citizen: Civic Engagement in a Mediatized Culture*, ed. René Glas et al. (Amsterdam University Press, 2019), 220.

↑32 Shannon Mattern, *A City is Not a Computer: Other Urban Intelligences* (Princeton University Press, 2021), "Contribute," OpenML, accessed September 1, 2025, https://www.openml.org/contribute, "Hugging Face," Hugging Face, accessed September 1, 2025, https://huggingface.co/huggingface.

↑34 Julian Kücklich, "Precarious Playbour: Modders and the Digital Games Industry," *The Fibreculture Journal* 5 (2005), https://five.fibreculturejournal.org/fcj-025-precarious-playbour-modders-and-the-digital-games-industry/.

↑35, ↑50 Kücklich, "Precarious Playbour".

↑36 Ethayarajh and Jurafsky, "Utility is in the Eye of the User: A Critique of NLP Leaderboards".

↑37 "Open LLM Leaderboard," Hugging Face, accessed September 1, 2025, https://huggingface.co/spaces/open-llm-leaderboard, "Leaderboard Overview," LMArena, accessed September 1, 2025, https://lmarena.ai/leaderboard, "Rankings," OpenRouter, accessed September 1, 2025, https://openrouter.ai/rankings.

↑38 Erica Klarreich, "A Decades-Old Computer Science Puzzle Was Solved in Two Pages," August 4, 2019, https://www.wired.com/story/a-decades-old-computer-science-puzzle-was-solved-in-two-pages/.

↑39 Sam Hind and Tatjana Seitz, "Cynical Technical Practice: From AI to APIs," *Convergence: The International Journal of Research into New Media Technologies* 30, no. 1 (2022): 29–48, https://doi.org/10.1177/13548565221133248.

↑40 David Ribes et al., "The Logic of Domains," *Social Studies of Science* 49, no. 3 (2019): 281, https://doi.org/10.1177/0306312719849709.

↑41 Ribes et al., "The Logic of Domains," 285.

↑42 Adrian Mackenzie, *Machine Learners: Archaeology of a Data Practice* (MIT Press, 2017).

↑43 Hind et al., "Challenges as Catalysts," 5.

↑44 "The State of Machine Learning Competitions: 2024 Edition," ML Contests, accessed September 1, 2025, https://mlcontests.com/state-of-machine-learning-competitions-2024/.

↑46 Thanks to Markus Stauff for posing this question during an invited talk in the ASCA Cities seminar series on 'Playable Cities' at the University of Amsterdam (NL) in March 2025.

↑47 To be eligible to compete in the Waymo Open Dataset Challenges in 2023, entrants, for instance, could not be a 'resident of Brazil, Quebec, Italy, or any country, state, province or territory subject to comprehensive OFAC [US Office of Foreign Assets Control] sanctions, including Cuba, Iran, North Korea, Syria, or the regions of Crimea, Donetsk or Luhansk of Ukraine,' "WOD Challenges Official Rules," Waymo, accessed May 1, 2025, https://waymo.com/open/terms/.

↑48 "Free Cloud Features and Trial Offer," Google Cloud, accessed May 1, 2025, https://cloud.google.com/free/docs/free-cloud-features.

↑**49** Jess Weatherbed and Lauren Felner, “Google Search Charged with Breaking EU Antitrust Rules,” *The Verge*, March 19, 2025, https://www.theverge.com/news/618168/google-search-eu-dma-charge-violation.

↑**51** Goree et al., “Attention is All They Need”.

↑**52** Everingham et al., “The PASCAL Visual Object Classes Challenge”.