Brains You Own, Connected: A Research Agenda for Community-Built Collective Decision Agents

CONNECT AI LAB
South Korea

Authors (138). First author, then every researcher who listed their name on the course's public research wall, in joining order.

Wonseok Jung · Genmon · Minjung Lee · myeongHui Jeon · SEUNGJUN-LEE · Daeseong Cheon · Jeonghan Lee · Iann Kim · Sihyun Kim · DaeHyeon Kim · sangyong JU · Seyun Kim · Sangryong Park · Kwangseob Kim · Unhee Yoo · KIKEUN NOH · Seung Hyun Kim · Hunjin Ju · soo ok Kim · SeungKwon Choi · Jihyun Ryu · SunYoung Kim · Insuk Choi · Inyoung Choi · EunHee Lee · HyeongSeok Kim · Incheol Jeong · Donghwi Kim · Sujin Kim · BO YEON KIM · cosmos · EUNSU CHO · Jungsuk Hwang · HyunJoo Cho · Eunyoung Lee · JIWOONG KANG · Junho Jeon · Sungwoon Hwang · Jongseok Lee · Kyoungsook LEE · Mihyun Lee · JaeSeok Kim · Hyoung Yeob Kim · Eugenia Kim · MinJin Kim · Sanghyeon An · Yeseul Jin · gyonghun · Jaejoong Kim · Songyi Kim · Il An · Chan Kyu Leo Park · Young Min Kim · Jiyoon Lee · Youngjun Kim · TaeUng Pae · NamSuk Kim · Sungpil Moon · Woncheol K · Hyunsuk Lee · JeongEun Lee · Teajung Kim · haup Jeong · Haeran Suh · Minju Lee · JANGJO YOON · Jiwoo Kim · Kyoung Soo Lee · Jinhyung Jo · Youngman Lee · Cheolwoo Kim · CHRIS · Jooeun Kim · Goheun Lee · JINSUNG LEE · Kyungbeom Kim · Yeri Lee · Joonho JANG · Baekahm Kim · Siwoo Park · Ho jin Youn · Junhyung Lee · YoungHo Kang · JeHwan Lee · hyokwionma · Kichul Pyeon · Kyoung Hwa Hyun · seongwook lee · Areyoung Kim · BongGyun-Kim · KI YOUNG LEE · Sunghoon Woo · SHINKUMJI · SUIN KIM · gilho shin · DONGSIK MYOUNG · Seongmin Lee · SOO YOUN HONG · gyeongseop lee · SONG JUNG SUK · jikja.kim · Eunjeong Kim · Minyoung Kim · Yunju Mork · miseon lee · Tae-Yoo Kim · Minkyoung Shin · EunAe Kim · SOOHYUNG KIM · 난성냥 · SANGHYUN BYUN · Julia Kim · SUKKYOUNG KIM · Kyoungsim Ahn · Kwangseob Kim · Yeonsik Jung · EUNJI PARK · 변상길 · 정 · 김경환 (Kyung Hwan Kim) · eunyoung kim · WeonJeong Shin · MINJU KIM · YONGJOO LEE · eunsa shim · WonSun NA · Terry IM · CHUNWEOL NHO · Jaby · Maeng Seong Un · JUHEE KIM · 장대규 (NST공법) · RIyun Kim · Jinnam Kim · KYUNGHA YONG · Sunghoon Park · James Yun · Dain_AI

Abstract

Decision-making AI is increasingly built as one large model owned by one organisation. We argue for a different substrate: many small decision policies ("brains"), each trained by an individual, published under a shared observation–action contract, and connected on demand for situations none of them was trained for. This turns questions that ensembles, model merging and multi-task reinforcement learning have studied inside single labs into open multiagent problems: interoperability, scoring strangers' contributions, aggregating disagreeing values, preserving diversity, and deciding when to verify or defer. We ground the agenda in a community pilot in which 144 participants of an online course uploaded 347 models for a company simulator and a robot world, then proposed their own connection rules. Averaging all 207 company brains beat the median but fell 25% short of the leaderboard leader. On unseen companies, a participant's rule (re-audition on the new company, then let the best brain run each department) beat every individual brain in three of four cases. On an unseen robot course, averaging collapsed to 0% success, while verifying by simulation only when brains disagree recovered 66%. We distil the proposals into testable mechanisms and outline the benchmarks, protocols and incentives that crowd-built brains need.

Keywords

Multiagent learning; collective decision-making; ensembles; policy aggregation; social choice; open model ecosystems; reinforcement learning

1 Introduction

A decision policy for running a small company, pricing a hotel or steering a robot can be tiny: a linear map from a few dozen observations to a handful of actions. What makes it valuable is not its size but what it encodes: the simulator its owner built, the reward weights they chose, and the experience they collected. Today such knowledge is either discarded or absorbed into one large model owned by one organisation. We ask what happens if it is instead kept as many small, *owned* policies that anyone can publish and anyone can *connect*.

We call this an *ecosystem of owned, connectable brains*. Each brain is trained by one person and published under a shared contract (the meaning and order of observations and actions) together with a record of where and how it was trained. When a new situation arrives, such as a company with different customers or a course the robot has never seen, a *connection rule* selects, weighs, routes or verifies the published brains to act. Ownership is preserved: a brain can be credited, licensed, revoked or improved by its owner without retraining anything else.

This picture sits between several well-studied ideas without being any of them. Ensembles [2] and weight averaging [7, 20] combine models, but assume members trained by one party on one objective. Mixtures of experts [8, 17] learn a router jointly with the experts. Policy distillation and Distral [14, 19] merge skills into one network, erasing ownership. Federated learning [9] shares gradients toward one global model. Ad hoc teamwork and zero-shot coordination [6, 18] connect agents that were trained apart, but for joint action among them rather than for one decision drawn from many advisers. Computational social choice [1] aggregates preferences, but voters are rarely policies whose competence varies by situation. An open ecosystem combines all of these difficulties at once, under adversarial and economic pressures that a single lab never faces, in the spirit of building models like open-source software [11].

The idea is blue sky because the ecosystem does not exist yet at scale. It is not fantasy, because we ran a small version of it. Section 3 reports what 144 course participants and 347 public models taught us: naive connection is easy and mediocre, the crowd's competence is concentrated and its diversity is thin, and the participants' own proposals point to mechanisms that already beat every

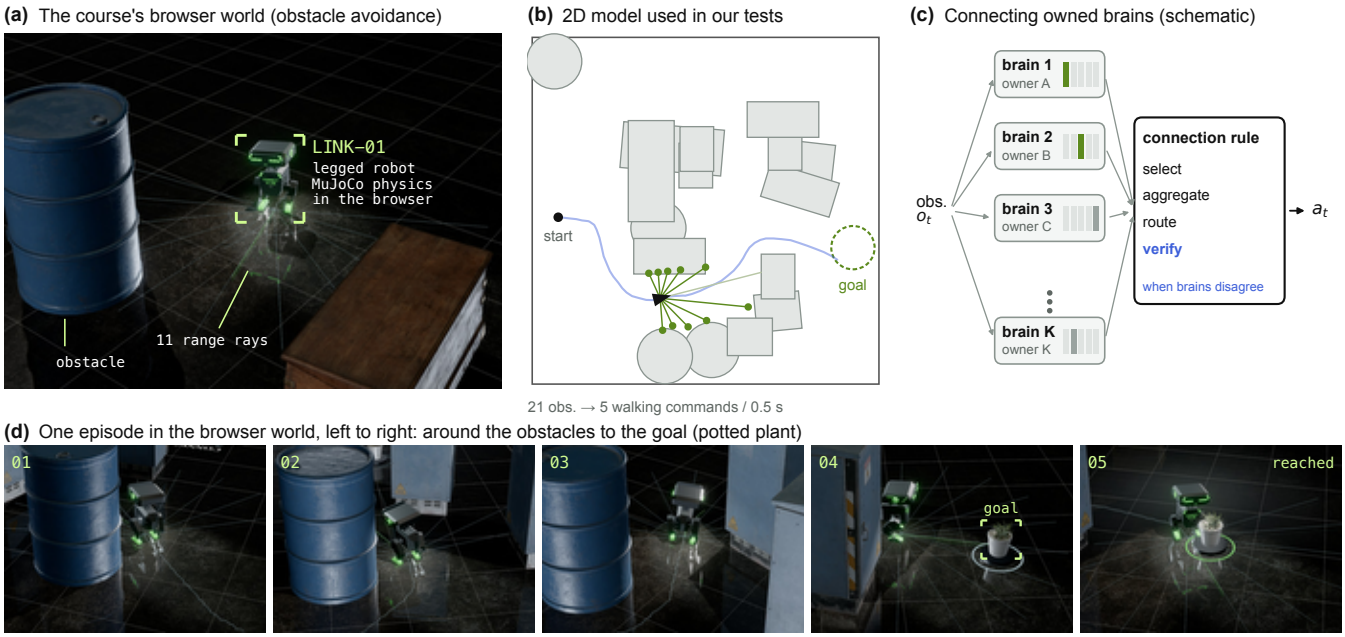


Figure 1: The robot world of the community pilot. (a) The course's browser simulator, in which participants trained obstacle-avoidance brains for a small legged robot, LINK-01 (MuJoCo physics in the browser; render captured from the public simulator). (b) The 2D walking model, calibrated to the browser physics, that we use for every robot result in this paper: course B, the path of the rule "verify when brains disagree" (blue) and the robot's 11 range rays at a step where the connected brains disagreed. (c) Schematic of connecting owned brains: each brain proposes a walking command and a connection rule turns the proposals into one action. (d) One episode of the simulator's built-in example policy, left to right.

individual in some settings. Section 4 turns those lessons into open problems.

2 Owned, connectable brains

Contract. A brain is useful to others only if its inputs and outputs mean the same thing everywhere. In our pilot each simulator fixes a versioned contract, for example six departments with named observations and named actions, and every published file states the contract, the simulator configuration it was trained on and a fingerprint of the course or company. Contracts make connection a lookup rather than an alignment problem, and make incompatibility detectable.

Ownership and provenance. Brains are published in personal repositories on a public model hub under their owner's name, with the reward weights, training curve and official test score. Provenance matters: it is what lets a connection rule ask *where* a brain was competent, not only *how* competent it claimed to be.

Connection rules. A connection rule maps the current observation and the set of published brains to one action. We distinguish four families. *Selection* picks brains (the leaderboard leader, the top k, or the winner of a fresh audition on the new situation). *Aggregation* combines their outputs (probability averaging, plurality or weighted votes, a veto). *Routing* assigns different brains to different roles or states, such as one brain per department. *Verification* checks a proposed action by short simulated rollouts or by a shared memory of past failures, and defers when the evidence is weak. Rules can be stacked, and each has a cost: auditions and rollouts consume simulator time.

Biological precedent. Honeybee swarms choose a nest site by quorum among scouts that have *inspected* candidates rather than by headcount [15]; ants mark failures and successes in a shared medium [3]; flocks coordinate through local rules [12]. Diverse groups can outperform groups of individually better solvers [5]. These precedents inspired many of the participant proposals below; they are inspiration, not evidence.

3 A community pilot

Setting. Over four weeks, participants of a public online course on reinforcement learning (CONNECT AI LAB) trained brains in two browser-based simulators and published them to the Hugging Face model hub. By November 2026, 144 distinct accounts had uploaded 347 models: 209 *company* brains and 138 *robot* models (all listed in Appendix A). The company simulator runs a month of an online business through six departments (staffing, inventory, marketing, customer service, refunds and security); a brain is a set of linear softmax heads, one per department, and is scored by monthly profit. Its default "general" company is calibrated to public datasets. The robot world asks a simulated legged robot to reach a goal while avoiding obstacles (Figure 1); a brain is a small MLP over 21 observations (goal distance and direction, 11 range rays, current speed and turn rate, and its previous command) choosing

among five walking commands, and we evaluate it in a 2D walking model calibrated to the browser's physics. Participants then posted, under course videos, their own proposals for how brains should be connected; we collected 318 comments under seven course videos, 64 of which proposed a connection rule or reported a result. The models are public; the comments were posted under course videos, some visible only to course members, and we report them only in aggregate or as anonymised paraphrases. We did not retrain any participant brain. Selection always uses held-out validation episodes, and test episodes are never used to choose a rule.

Lesson 1: naive connection beats the median, not the leader. We re-scored all 207 loadable company brains on 40 held-out test months (Table 1, Figure 2). Averaging all brains' action probabilities earned 1171 versus a median brain's 1115, but the brain that led the validation leaderboard earned 1568: the average fell 25% short of it. The same pattern held on four unseen company variants built from other industries' demand, refund and fraud profiles: averaging all brains beat the median by 2–6% and lost to the leader by 8–33%. Small committees of the top five brains came within 3% of the leader and, on the hotel variant, beat every individual brain.

Lesson 2: diversity is thin and competence is concentrated. 175 of the 207 company brains (85%) kept the default reward weights. Grouping brains by their actions on shared probe states gives 128 behavioural classes, but the two largest hold 71 brains (34%), and 119 brains form classes of their own (Figure 3). Three robot uploads from two accounts were byte-identical. The strongest company brain came from one participant who published nine successive versions and changed the reward weights; custom weights appear in 8 of the 20 validation leaders but in only 15% of all brains. Averaging a crowd of near-copies cannot add information, and it dilutes the few brains that have some.

Lesson 3: participants' rules work when they add information. One participant proposed to *"follow the best brain for each part, then test again and follow whoever scores highest"*. We implemented this as a per-department team re-fitted on 12 audition months of the new company (disjoint from test months), choosing among the ten validation leaders. It beat every individual brain on the retail, hotel and risky variants (Table 1); on the lecture variant, where the leaderboard leader ranked only in the top 17%, a plain re-audition of all brains recovered the best one. The common thread is that these rules spend a little experience in the new situation before trusting anyone.

Lesson 4: in the physical world, bad aggregation is dangerous and verification pays. In the robot world we connected the six distinct obstacle-avoidance brains on courses that were *not* in their training sets, over 100 episodes per course (Figure 4). On the hardest unseen course, averaging all brains collided in every episode and never reached the goal: most brains had never seen such a layout and outweighed the one that could handle it. Participants proposed weighting votes by confidence or by public score, giving any brain a veto in front of danger, sharing a "wrong-answer notebook" of situations that preceded collisions, and verifying by simulation only when brains disagree. Quorum-gated verification, which tries each disputed action a few steps ahead in copies of the simulator with fresh walking noise, raised success from 0% to 66%, at the price of about 48 simulated steps per real step on that course. The minority veto, perhaps the most intuitive safety rule, was the worst: when most voters are incompetent, every useful action is vetoed by someone. Here the public score was a good guide: the brain with the highest self-reported score, trained on none of these courses, matched or beat every connection rule, and both overriding it by majority and filtering it through the shared failure memory left it no better. On two further, easier courses every rule succeeded in at least 93% of episodes.

Table 1: Company simulator: mean monthly profit on 40 held-out test months. "General" is the company the brains were trained on; the other four are unseen variants. Selection uses validation months of the general company, or 12 audition months of the new company where marked. Bold: beats every one of the 207 individual brains.

	General	Lecture	Retail	Hotel	Risky
Median single brain	1115	1440	896	1527	1015
Best single (oracle, test)	1568	1719	1396	1993	1450
Leaderboard leader	1568	1633	1396	1973	1450
Average, all 207	1171	1501	933	1551	1080
Trust-weighted average, all	1192	1512	961	1593	1114
Average, top 5	1528	1598	1354	**1997**	1411
Per-department team, top 10	1560	1609	1377	1958	1426
Re-audition, best of all	–	1719	1396	1993	1450
Per-dept. team re-fitted	–	1653	**1434**	**2021**	**1461**

Lesson 5: small is not a compromise. Most robot uploads were 0.8B-parameter language models (LoRA adapters on Qwen3.5-0.8B in 133 of the 138 robot repositories) fine-tuned to imitate a small MLP brain (usually the participant's own), then refined with reinforcement learning. The course software logged teacher and student on 20 closed-loop episodes at upload time for 136 such repositories from 71 accounts (Table 2); 126 of them were still public at our snapshot. On obstacle avoidance the small teacher MLPs succeeded in 87% of episodes and their language-model students in 56%, although students agreed with teachers on 69–72% of logged decisions. On route following and object choice the gap almost closed (98% vs. 92%, 85% vs. 85%), but on goal reaching it stayed (81% vs. 58%). This is covariate shift [13] in miniature, and one participant diagnosed it as such in their own write-up. A connectable ecosystem of small, cheap, inspectable brains is not a fallback from large models; for closed-loop control it can be the better default.

Lesson 6: connection is also a good place to start learning. We also used connected brains as the starting point for further training on each unseen company: 30 PPO iterations of 24 simulated months, three seeds each (Figure 5). Starting from the leaderboard leader or from the weight average of the top 10 gave the highest validation profit at iteration 0 in all four companies, and further training lowered it slightly in every case (hotel: 1999 to 1930 and 2008 to 1929). Training from scratch was still far below after 30 iterations (hotel: 221), and imitating the hand-written rules first reached 1492. Averaging the weights of all 207 brains started below the top-10 average in every company: here too, the crowd helps when it is filtered.

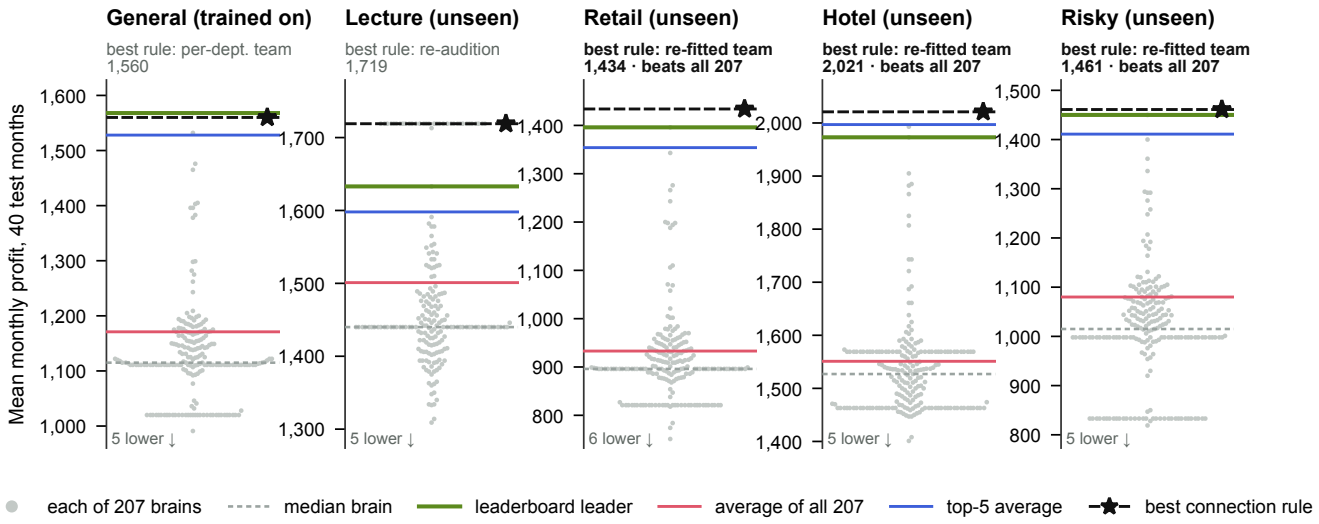


Figure 2: Company simulator, 40 held-out test months. Each grey dot is one of the 207 participant brains (the few lowest are cut off and counted at the bottom); lines mark the median brain, the validation leaderboard leader, the average of all 207 brains, the top-5 average and the best connection rule in each company. On three of the four unseen companies a participant's rule (re-fit a per-department team on 12 audition months) beats every individual brain.

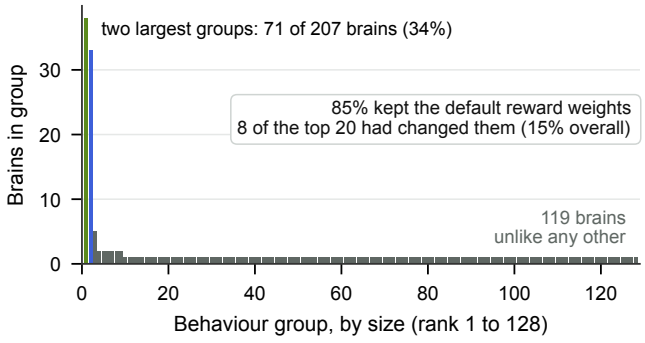


Figure 3: Behavioural diversity of the 207 company brains: 128 groups of brains that choose the same actions on shared probe states, ordered by size.

Table 2: Robot world: small MLP teachers and their 0.8B language-model students, as logged by the course software at upload (20 closed-loop episodes per repository). *n*: repositories. Success rate of the teacher, of the student after supervised imitation (SFT) and after RL refinement; agreement: share of logged decisions in which the student chose the teacher's command.

		Success			Agreement	
Task	*n*	Teacher	SFT	RL	SFT	RL
Obstacle avoidance	51	0.87	0.53	0.56	0.72	0.69
Goal reaching	40	0.81	0.56	0.58	0.78	0.76
Object choice	23	0.85	0.82	0.85	0.98	0.97
Route following	22	0.98	0.89	0.92	0.80	0.79

4 Research agenda

Contracts that evolve. Simulators change; the robot world already has four observation layouts across its tasks, two of them for the same task. We need versioned contracts with adapters and with tests that detect when two brains that look compatible are not.

Trust without a single leaderboard. A leaderboard is one company, one course and one seed set. The leader dropped from the top 0.5% to the top 17% on one unseen variant, and auditions fixed this at a measurable cost. Open questions: how many audition episodes a rule should spend, how to share audition results among users without leaking test sets, and how to resist owners who tune brains to the leaderboard.

Aggregation under disagreeing values. Brains trained with different reward weights disagree for good reasons. Social choice offers axioms [1], but its voters are not policies whose competence varies by state. Condorcet-style guarantees [4] assume independent, better-than-chance voters; our pilot violates both. What are the right axioms when voters are near-copies and competence is local?

Paying for diversity. If 85% of contributors keep the defaults, the ecosystem must reward contributions that add *marginal* information to connected decisions, not standalone score. Credit assignment over coalitions [16] is a natural start; the commons literature [10] warns that incentives, not algorithms, decide whether a shared resource survives.

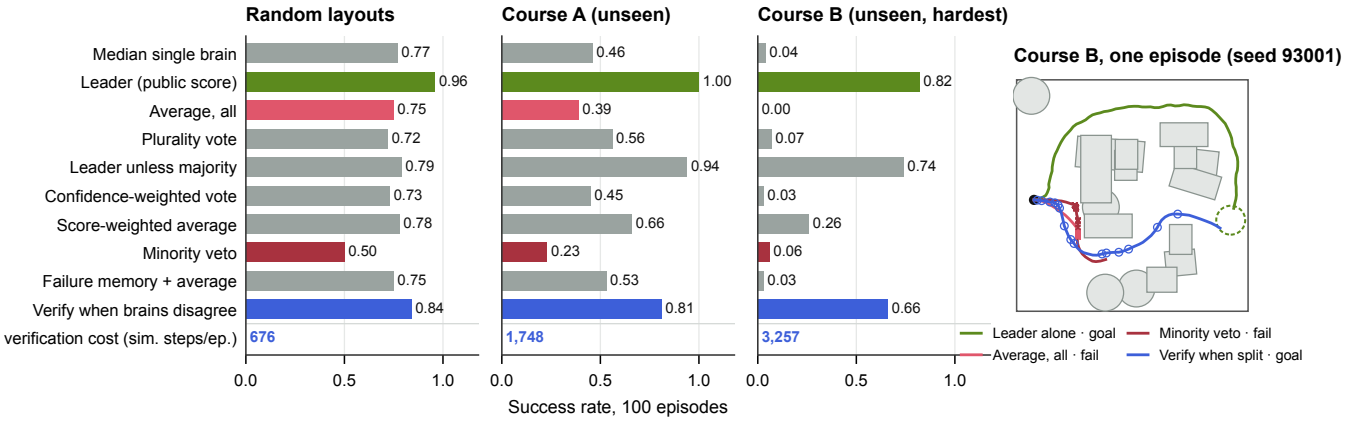


Figure 4: Robot world, courses absent from the connected brains' training sets (five brains each). Left: success rate over 100 episodes for ten ways of acting; the bottom row is the simulated steps per episode that verification spent. Right: one episode on course B with the same seed for four rules (crosses: collisions; circles: steps where brains disagreed and verification ran).

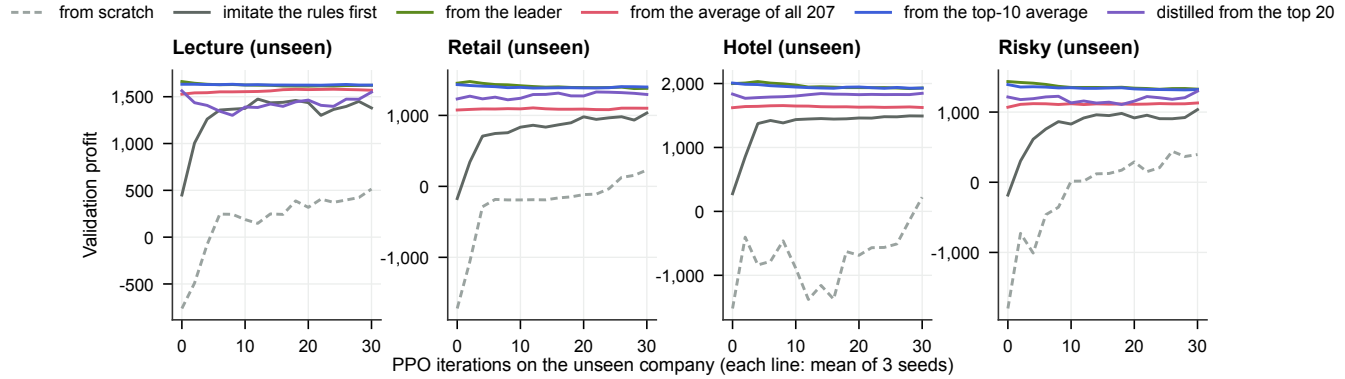


Figure 5: Further training on the four unseen companies from six starting points: validation profit over 30 PPO iterations (24 simulated months each), mean of three seeds.

Verification and deferral. Quorum-gated verification turned a catastrophic average into a usable controller, but it used the simulator itself as a world model. Replacing it with learned world models, sharing failure memories across owners as ants share pheromone [3], and deciding when to defer to a human are open, safety-critical problems.

Ownership and safety. Owned brains can be revoked, licensed and attributed, but can also be poisoned or duplicated. We need provenance checks, duplicate detection (which already mattered in our pilot) and rules that remain safe when some contributors are adversarial.

5 Outlook

We have argued that decision AI could be grown as an ecosystem of many small, owned brains rather than built as one large model, and that the multiagent community is best placed to make that ecosystem reliable. Our pilot is small, its simulators are simplified, and its participants were students of one course; none of its numbers is a benchmark result. What it shows is that the questions are real, measurable and already within reach of a few hundred people with laptops. We will release the scripts, the list of public models (Appendix A) and the anonymised idea comments, so that others can reproduce the pilot and propose better connection rules.

Use of AI tools. We used Claude Opus 5.5 (Anthropic), through the Claude Code agent, to write and run the evaluation scripts, collect and summarise the course-video comments, and draft and edit this text. The connection rules evaluated were proposed by human participants; every reported number is produced by the scripts in the supplementary material, which also lists the prompts.

References

[1] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice.* Cambridge University Press, Cambridge, UK. doi:10.1017/CBO9781107446984

[2] Thomas G. Dietterich. 2000. Ensemble Methods in Machine Learning. In *Multiple Classifier Systems: First International Workshop, MCS 2000, Cagliari, Italy, June 21–23, 2000, Proceedings (Lecture Notes in Computer Science, Vol. 1857)*, Josef Kittler and Fabio Roli (Eds.). Springer, Berlin, Heidelberg, 1–15. doi:10.1007/3-540-45014-9_1

[3] Marco Dorigo, Vittorio Maniezzo, and Alberto Colorni. 1996. Ant System: Optimization by a Colony of Cooperating Agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 26, 1 (1996), 29–41. doi:10.1109/3477.484436

[4] Bernard Grofman, Guillermo Owen, and Scott L. Feld. 1983. Thirteen Theorems in Search of the Truth. *Theory and Decision* 15, 3 (1983), 261–278. doi:10.1007/BF00125672

[5] Lu Hong and Scott E. Page. 2004. Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proceedings of the National Academy of Sciences* 101, 46 (2004), 16385–16389. doi:10.1073/pnas.0403723101

[6] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. 2020. "Other-Play" for Zero-Shot Coordination. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020) (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 4399–4410. https://proceedings.mlr.press/v119/hu20a.html

[7] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. 2023. Editing Models with Task Arithmetic. In *The Eleventh International Conference on Learning Representations (ICLR 2023)*. https://openreview.net/forum?id=6t0Kwf8-jrj

[8] Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. 1991. Adaptive Mixtures of Local Experts. *Neural Computation* 3, 1 (1991), 79–87. doi:10.1162/neco.1991.3.1.79

[9] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS 2017) (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Jerry Zhu (Eds.). PMLR, 1273–1282. https://proceedings.mlr.press/v54/mcmahan17a.html

[10] Elinor Ostrom. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, UK. doi:10.1017/CBO9780511807763

[11] Colin Raffel. 2023. Building Machine Learning Models Like Open Source Software. *Commun. ACM* 66, 2 (2023), 38–40. doi:10.1145/3545111

[12] Craig W. Reynolds. 1987. Flocks, Herds and Schools: A Distributed Behavioral Model. *ACM SIGGRAPH Computer Graphics* 21, 4 (1987), 25–34. Proceedings of SIGGRAPH '87. doi:10.1145/37402.37406

[13] Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. 2011. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS 2011) (Proceedings of Machine Learning Research, Vol. 15)*, Geoffrey Gordon, David Dunson, and Miroslav Dudík (Eds.). PMLR, Fort Lauderdale, FL, USA, 627–635. https://proceedings.mlr.press/v15/ross11a.html

[14] Andrei A. Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. 2016. Policy Distillation. In *4th International Conference on Learning Representations (ICLR 2016)*. San Juan, Puerto Rico. https://arxiv.org/abs/1511.06295

[15] Thomas D. Seeley. 2010. *Honeybee Democracy*. Princeton University Press, Princeton, NJ. https://press.princeton.edu/books/hardcover/9780691147215/honeybee-democracy

[16] Lloyd S. Shapley. 1953. A Value for *n*-Person Games. In *Contributions to the Theory of Games, Volume II*, Harold W. Kuhn and Albert W. Tucker (Eds.). Number 28 in Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 307–317. doi:10.1515/9781400881970-018

[17] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *5th International Conference on Learning Representations (ICLR 2017)*. https://openreview.net/forum?id=B1ckMDqlg

[18] Peter Stone, Gal A. Kaminka, Sarit Kraus, and Jeffrey S. Rosenschein. 2010. Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2010)*, Vol. 24. AAAI Press, Atlanta, Georgia, USA, 1504–1509. doi:10.1609/aaai.v24i1.7529

[19] Yee Whye Teh, Victor Bapst, Wojciech Marian Czarnecki, John Quan, James Kirkpatrick, Raia Hadsell, Nicolas Heess, and Razvan Pascanu. 2017. Distral: Robust Multitask Reinforcement Learning. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc., 4496–4506. https://proceedings.neurips.cc/paper_files/paper/2017/hash/0abdc563a06105aee3c6136871c9f4d1-Abstract.html

[20] Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. 2022. Model Soups: Averaging Weights of Multiple Fine-Tuned Models Improves Accuracy without Increasing Inference Time. In *Proceedings of the 39th International Conference on Machine Learning (ICML 2022) (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 23965–23998. https://proceedings.mlr.press/v162/wortsman22a.html

A Public models

All 347 repositories on the Hugging Face hub that we evaluated, from 144 accounts, as of November 2026. Each line is one account (`owner /`) followed by its repositories. Every repository is named `owner/connectailab-rl-name-teacher` (company simulator) or `owner/connectailab-rl-name-link01` (robot world), and we print only *name*; names marked † do not follow this pattern and are printed in full. 133 of the 138 robot repositories contain a LoRA adapter on Qwen3.5-0.8B.

Company-simulator brains (209 repositories, 132 accounts)

abracsus / ewbrain-v1, ewbrain-v2
aha123456789 / connectailab-rl-changer-v1
Aiden-ing / aiden-v6
AlexLEEeeee / alvis-v1, alvis-v2, alvis-v3, alvis-v4
allapp / jay-v1, mong-v1
Aoife1111 / aoife-v1, aoife-v9
appssy3 / solo-build-h-v2
artnfull / artnfull-v1
attitudemetal / aion-v1, aion-v2
avoca8Hug / avoca-v2
bgkim3514 / connectailab-rl-km1-v1-teacher-v1, connectailab-rl-km1-v1-v1, connectailab-rl-km1-v1.1-v1
bmo0206 / jay-v1, myeongok-v1
Bomdeang / maengseongun-v1
boyeon-connectai2026 / boyeon-v1
byeon120 / connectailab-rl-lifeenergy-v3-v3
casareborgia / leeseungjun-v1
Charlottej23 / decider-v1, decider-v2, decider-v3
coldpink / sharala-v1
connect-ai-lab-aimo / jay-v1, juju-v1
dadajikgu / dada-v1
Daeseong73 / daeseongcheon-v1
dain-studio / yggdrasil-v1
DaOn-ai / connectailab-rl-daonai-v1, daonai-v1, daonai-v2
devorah-ai-2016 / devorah-v1
dni-global-healthcare / connectailab-rl-dni-v1-v1, dni-v3
eplus70 / eplus-v1
euna88 / nextforest-v2
eunsasim / eunsa-v1
fanglinxun / aicompanydecision-v1
GAON-PAPA / chris-v4
glassidenttrab / jay-v1
goodbaat / goodbaat-v1
Guccimam / shibalove-v2, shibalove-v34
gutgai / sirius-v2
happyappletree / apple-v1
hiupjump / harugabrain-v2, harugabrain-v3
hope6803 / jarample-v1
huggsook / saffobrain-v1
hwihwalab / CHOCO-BRAIN_GUARD-v1
ianrua / jrua-v5
ic4u2u / yong-v7
ieum3355 / ieum-v1
imleadingmylife / jay-v1
INFJ-FATHER / salus-v1
insukchoi / mochi-v1
IRELAB / irelab-v1
irobtec / irobtecbrain-v1
izblue / izblue-v1, izblue-v2, izblue-v4
James-1971-source / connectailab-rl-jamesyun-v1
jangjoYoon / justinbrain-v1, justinbrain-v2, justinbrain-v7
Janovation / nst-v1
Jasmin71 / jasmin-v7
jb2jb2kr / jay-v1
jchann00 / jay-v1, jcan-v2
jfcrown7 / eskim-v2, jay-v1
joaas / jay-v1, joaas-v1, kkommii-v1
johnkim7788 / kebi-v1
jongchullee / jongchul-v1, leejongchul-v2
jsy6280 / connectailab-rl-solozen-v1, jay-v1
kareyoung09 / ideun-v1, ideun-v2
key1739 / connectailab-rl-coco-brain-v1
Kimsihyun / shdm-v1
kiodo / brena-v2, jay-v1
kiso2 / jmbrain-v1
kiwi2030 / kiwi-v2
kmedicos / youngsuk-v1
kounghwa / jay-v1, kyounghwa-v10
kwanghyunjoo / kwanghyunjoo-v1, kwanghyunjoo-v2
kymbms / kodari-v1
kyoungsook70 / tbrain-v1
lazysage / lazybrain-v1
leegoheun / duking-v1
leepori / connectailab-rl-poribrain-v1
LeeSanghyun / LeeSanghyun-v1, LeeSanghyun-v13, LeeSanghyun-v15, LeeSanghyun-v2, LeeSanghyun-v3, LeeSanghyun-v4, LeeSanghyun-v5, LeeSanghyun-v6, LeeSanghyun-v7
litzor / litzor-v1, litzor-v2, litzor-v3
lookarooka / novice-v1, novice-v2, novice-v3
lucaskim72 / lucasbrain-v3
lucyhome / lucy-v1
Luna002-Luna75 / connectailab-rl-lunaluna-v1
Lunation / lunar-v1
maro-aigent / jay-v1, maro-v1
Martinyooo / martinyooo-v1
mehwa502 / connectailab-rl-navillera-v1, mehwa502-v1, navillera-v1
meonghui / yacht-v1
Mike9 / connectailab-rl-oasisbrain-v1-v1
minju2026 / milkbrain-v1
minjunglee27 / highj, highj-v2
momolight / momonana-v1
moneyclubx / connectailab-rl-kimchulmin-v1-v1
MOONSIK89 / jay-v1
nawonsun / connectailab-rl-lalabrain-v1-v1, jay-v1, lalabrain-v2
netsesang / jay-v1, netsesang-v1, netsesang-v2
nexflow / cocobrain-v1, cocobrain-v3
ohyou / connectailab-rl-oddissey-v1-v1, jay-v1
onethingtree / onetbrain-v1, onetbrain-v2, onetbrain-v3
piramid777 / jay-v1
pjw1356 / realnote-v1
pluskkm79 / jay-v1, pluskkm79-v1
pwcosmos / pwcosmos-v1, pwcosmos-v3
rhemslkim / jay-v1
seyunprobiz / seyunbrain-v1
SHSemi / seung-hyun-kim-v1
Siwooo / jay-v1
Slowtownstory / apple-v1
sm-routinefit / slabrain-v1
sol-5 / cassbrain-v1
sotdanji / connectailab-rl-sodu-v1, jay-v1
spmoon / jay-v1
spoindo / spoindo-v1
Spyfood / ragu-v1
station124512 / connectailab-rl-goodbong2-v1-v1
sung-hoon / connectailab-rl-kribrain-v1
sunyoungKim / lco-v1, lco-v2, lco-v3, lco-v4, lco-v5
suryeon123 / connectailab-qwen-soo-v1-teacher†, soo-v1
sxixkxx / sxixkxx-v1
tazza11 / simba-v2, simba-v3, simba-v4
the-eschatos / shp-brain-v3, SHP-v1, shp-v2, shp-v3
THEALPS / kwan-ok-hur-v1
tlsrmawl / moa-v1
TUpreneur / tupreneur-v1
WeonJeong / maru-v1, maru-v2
whcho33 / joy-v1
white100big / jay-v1
WithHaebom / haebom-v1
WonseokJayJung / beerbrain-v1, jay-v1
wsy-1003 / soribrain-v1, soribrain-v2
yegom / jay-v1, yegomfloor-v2, yegomfloor-v4
Yeri1104 / jay-v1, yeri-v1
yonga76 / altos-v1
youn485 / connectailab-rl-jinbrain-v1, connectailab-rl-jinbrain-v4, connectailab-rl-jinbrain-v5
zero885814 / grid-empathy-v1
ZeroGravityLab / joonho-jang-v1

LINK-01 robot brains (138 repositories, 72 accounts)

abracsus / abracsus-v1, abracsus-v2, abracsus-v3, abracsus-v4
aha123456789 / changer-v1, connectailab-rl-changer-v1-link01-v1
allapp / jay-v1, mong-v1
appssy3 / solo-build-h-v1, solo-build-h-v2, solo-build-h-v3, solo-build-h-v4, solo-build-h-v5
arabellako22 / malibu-v1
artnfull / artnfull-v1, artnfull-v2, artnfull-v3, artnfull-v4
attitudemetal / aion-v1, aion-v2
avoca8Hug / avoca-v1
bigeye35 / healmu-v1
Bomdeang / maengseongun-avoid-v1, maengseongun-avoid-v2, maengseongun-avoid-v3
boyeon-connectai2026 / boyeon-kim-v1
caley-kim / namsukkim-v1
Charlottej23 / superwoman-1-v1, superwoman-2-v2
connect-ai-lab-aimo / juju-v1
creamapple / koco-v1
dadajikgu / dada-v1
dedalus2k14 / deda-v1
devorah-ai-2016 / devorah-v1-avoid-obstacle, devorah-v1-avoide-obstacle, devorah-v1-choose, devorah-v1-follow-route

euna88 / nextforest-v2
eunsasim / connectailab-rl-eunsa.objective-v1-link01-v1
Guccimam / shibalove-v1
gutgai / sirius-v1, sirius-v2
happyappletree / appletree-v1
hope6803 / jarample-v1, jarmample-v2, jarmample-v3
hwihwalab / CHOCO-BRAIN-v1
ianrua / lee-hyunju-v1
ieum3355 / ieum-v1
IRELAB / irelab-v1, irelab-v2-choose
irobtec / irobtecbrain-v1
izblue / izblue-v1
James-1971-source / jamesyun-v1
Jasmin71 / mihyun-v1
jb2jb2kr / cherrybrain-v1
jyungkang / jyungkang-v1
key1739 / connectailab-rl-coco-brain-v2
Kimsihyun / jay-v1, srdm-v1
kiodo / brena-v1
kiwi2030 / kiwi-v1
kyoungsim / kongi-v1
leegoheun / duking-v1, duking-v2, duking-v3, duking-v4
lookarooka / jay-v1
lucaskim72 / lucasbrain-v1
lucyhome / lucy-choose-v1, lucy-target-v1
Luna002-Luna75 / lunaluna-object-v1, lunaluna-obstacles-v1, lunaluna-route-v1, lunaluna-target-v1
Lunation / lunar-v1
maro-aigent / maro-v1
Martinyooo / martinyooo-v1
minju2026 / ducky-v1
minjunglee27 / highj-avoid-v1, highj-choose-v1, highj-follow-v1, highj-goto-v1
mogushop520 / moguverse-v1, moguverse-v2
NADANAYO / nadanayo-v1
nawonsun / connectailab-rl-robo-yuna-v1
nexflow / coco-v1, coco-v4, coco-v5, coco-v6
onethingtree / onetbrain-v1
PANGSE74 / aiot-v1
piramid777 / connectailab-rl-matop-link01-v2
pluskkm79 / jay-v1
rhemslkim / rere-v1, rere-v2, rere-v3, rere-v4, rere-v5
seyunprobiz / seyunduck-brain-v1, seyunduck-obstacles-v1, seyunduck-route-v1
SHSemi / connectailab-rl-seung-hyun-kim-v1-link01-v1, connectailab-rl-seung-hyun-kim-v1-link03-v1, jay-v1
spmoon / jjangu-v1, jjangu-v2, jjangu-v3, jjangu-v4, jjangu-v5
spoindo / spoindo-goto-v1, spoindo-v1
station124512 / jay-v1
tazza11 / tazza11-v2, tazza11-v3, tazza11-v4
tgc97057 / gs.lee-ana-v01-v1
the-eschatos / toty-v1, toty-v2
tskim81 / kimtaesun-avoid-obstacles-v1, kimtaesun-choose-object-v1, kimtaesun-follow-route-v1, kimtaesun-target-v1, kimtaesun-target-v1-link01
TUpreneur / tupreneur-object-v1, tupreneur-obstacles-v1, tupreneur-route-v1, tupreneur-route-v2, tupreneur-target-v1
WeonJeong / maru-v1
WithHaebom / withhaebom-v1
wsy-1003 / chamsori-goal-v1, chamsori-goal-v2
zero885814 / grid-empathy-v1, grid-empathy-v2, grid-empathy-v3, grid-empathy-v4