Assembloid Agency: Unreal Engine API for brain-on-a-chip platforms

Jenn Leung
University of the Arts London
London, United Kingdom
j.leung@fashion.arts.ac.uk

Chloe Loewith
University of Cambridge
Cambridge, United Kingdom
cl969@cantab.ac.uk

Abstract

Assembloid Agency proposes an open-source Unreal Engine API for interfacing with brain-on-a-chip platforms[1], where the API mediates between living neurons cultured on high-density microelectrode arrays and simulated game environments. Building on 'Organoid Array Computing: The Design Space of Organoid Intelligence' [12], which speculates on a future where three-dimensional brain cultures assemble into more complex cognitive infrastructures and hence may become increasingly 'designable' and 'playable' organisms[2]. We extend this design space to apply games design principles to context engineering for organoid intelligence, treating organoids as polycomputational agents trained through reinforcement learning[20].

Our proposed plugin exposes functions for stimulation, recording, visualization, and real-time control of neuronal cultures within Unreal Engine while providing various RL-based game templates. This allows researchers to rapidly prototype experimental contexts, whether using biological wetware, spiking neural networks, or EEG stand-ins. Game templates will be built on Unreal Engine's Learning Agents plugin to support single- and multi-organoid training scenarios. Theoretically, the API acts as a mediating membrane where biological and algorithmic agencies "contaminate" each other[18], dissolving model/reality divides and enabling distributed cognition across neurons, code, and simulated space[7]. We address risks such as value capture by incorporating reward-schema versioning, rotating reward schedules, and stress monitoring[15].

Assembloid Agency embraces the dual embodiment of biological intelligence. Here, "assembloid" refers to assembling organoids into new computational ecologies, while "agency" invokes not only the game engine as an experimental agent, but also the negotiated play of agency between biological and synthetic actors. By designing an API for game engines, we also anticipate the possibilities of interdisciplinary applications across games, interactive experiences, AI benchmarking, to architectural design, while foregrounding ethical and aesthetic guardrails organoid intelligence.

1 Unreal Engine API for brain-on-a-chip platforms

We propose developing an Unreal Engine API to interface with brain-on-a-chip platforms, which specifically refer to 3D cultures of brain organoids integrated with high-density microelectrode arrays capable of bidirectional electrical communication [1]. While originally developed for biomedical research, these platforms are now emerging as a novel substrate for biological computing, known as organoid intelligence (OI) [20]. Pioneering work in this field by Cortical Labs and researchers from Johns Hopkins University has demonstrated the potential of these OI systems[10, 20].

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Creative AI Track.

Building on recent breakthroughs, including FinalSpark's NeuroPlatform [5] and Cortical Labs' DishBrain project [10], we aim to extend traditional AI gameplay paradigms to incorporate living neural systems. Our proposed Unreal Engine plugin will provide tools to: 1) record neuronal activity in real-time 2) deliver targeted electrical stimulation 3) train living neurons within virtual environments through reinforcement learning game templates.

By creating an API that interfaces with biocomputers like the CL1 and simulators such as NEST simulator[6, 9], we enable a new generative approach to game design that incorporates living neural networks directly into game engines. This platform will allow researchers and developers to prototype organoid-based systems in rich, interactive 3D environments.

2 Methodology and related work

2.1 Organoids as organisms of polycomputation

Prior work by Leung et al. in 'Organoid Array Computing' establishes a framework for designing with Organoid Intelligence (OI), proposing that brain organoids perform multiple simultaneous computations through chemical, mechanical, and biological processes. These interconnected systems "form the basis of a polycomputational system" [12]. This perspective extends Bongard and Levin's paradigm of biological polycomputation. As Bongard and Levin observe: "Biology is rife with polycomputing at all scales." Often, a single substrate could perform multiple computational tasks simultaneously [3], for example, OI systems combine morphogenetic cultures, microelectrode arrays, vascular scaffolds, and machine learning algorithms to create evolvable polycomputational agents.

Meanwhile, OI system architectures often depend on computational approaches stacked together, such as physical reservoir computing, where the organoid itself functions as a neural network layer [3], as well as multiple layers of reinforcement learning, which has proven effective for training organoids [10]. As reinforcement learning operates recursively across all layers of an OI system, it is apparent that a multilayered approach is necessary to mirror the poly-computational nature of the biological substrate itself.

2.2 API as mediating membrane

The Assembloid Agency API is designed to be a bi-directional interface between the game engine and neuronal cultures. The game engine delivers closed-loop electrical stimulation to organoids via high-density microelectrode arrays based on dynamic gameplay feedback[10] , while organoids with neuroplastic spiking activities are mapped to reinforcement learning frameworks that reshape observations and reward functions[10, 5].

Following Parisi's framework, this membrane enables mutual "contamination" between biological and algorithmic agencies, dissolving conventional model/reality distinctions. The system, then, demands continuous renegotiation of fundamental reinforcement learning constructs: what constitutes state, action, or reward is dynamically co-defined.

Through the membrane between the organoid layer and the virtual environment, agential fluidity is also enabled by this interface. Following Hutchins' notion of distributed-cognition, cognition, much akin to Bongard and Levin's poly-computational framework, is present across artefacts, environments, and bodies[7]. With this, we also speculate that the API membrane facilitates the distribution of agency across scales.

2.3 Game Design as Agency Play

This approach draws on a tradition of using games as testbeds for artificial intelligence. Early examples such as Tesauro's TD-Gammon demonstrated that reinforcement learning through self-play could reach superhuman levels in backgammon without explicit strategic programming[21]. Two decades later, DeepMind's deep Q-network (DQN) achieved a major milestone by learning to play a diverse set of Atari 2600 games directly from raw pixel inputs, using a combination of convolutional neural networks, Q-learning, and experience replay [14]. The same architecture and hyperparameters were applied across multiple games, establishing the Arcade Learning Environment as a standard benchmark for general RL agents. AlphaZero combined deep neural networks with Monte Carlo Tree Search to learn and master chess, shogi, and Go entirely through self-play, starting from random play with only the basic rules provided, without any human game data or handcrafted domain knowledge [19]. OpenAI Five applied large-scale multi-agent RL to Dota 2, a partially observable, real-time strategy game requiring coordination, long-horizon planning, and adaptation to novel tactics [17].

Recently, large language models have been adapted to play games without explicit RL training, using in-context learning alone. In the Street Fighter III experiment[16], models received textual descriptions of the game state and recent moves, then selected the next move in real time. Interestingly, smaller, lower-latency models often outperformed larger ones, suggesting that reaction speed can outweigh raw model capacity in time-critical domains. These experiments also surfaced failure modes (hallucinating invalid actions, or refusing to play) that highlight the importance of grounding agents in the constraints of their environment.

Placing wetware into this lineage reframes these AI benchmarks: the "policy" is not a set of weights in silico but a living, plastic network whose synaptic configurations evolve in response to designed feedback loops. By providing standardized game templates, reward schemes, and data pipelines, the Unreal Engine API situates organoid intelligence within a comparative framework used for software agents. This continuity allows for direct methodological borrowing, from self-play to reward shaping, while extending the scope of what counts as an agent in game-based AI research.

3 Poly-computational layers of Assembloid Agency API

3.1 CL1/ NEST

The proposed toolkit will be implemented in two parallel configurations to accommodate both biological and simulated neural systems:

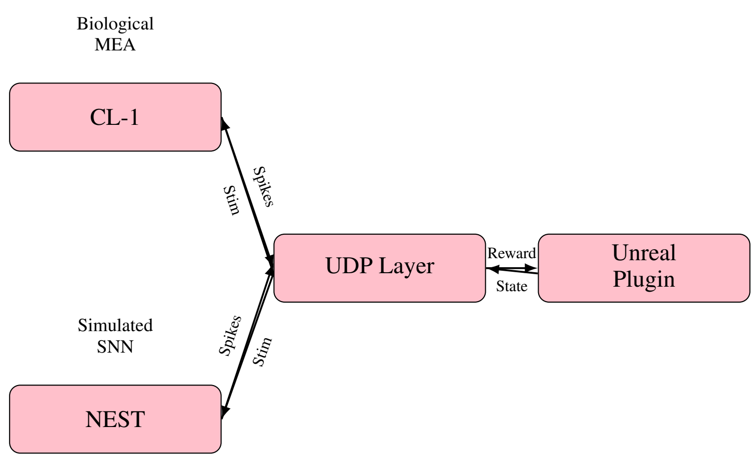

Figure 1: Compact bidirectional pipeline between biological (CL-1) or simulated (NEST) neural systems and Unreal Engine. Arrow pairs show closed-loop symmetry: neural spikes $\rightarrow$ engine state (top arrows) and engine rewards $\rightarrow$ neural stimulation (bottom arrows).

3.1.1 CL1 Integration

For biological computing, the plugin integrates with Cortical Labs' CL1, a neuromorphic platform where living neurons cultured on high-density microelectrode arrays process information in real time through closed-loop electrophysiology. CL1 combines optimized cell culture conditions, FPGA/ASIC hardware for millisecond-precision stimulation and recording, and a Python API (CL-API) for bidirectional control [9]. In their recent Nature Reviews Bioengineering article, Kagan proposes that commercially viable devices such as CL1 could be used to benchmark generalized intelligence abilities, such as navigating game environments.

3.1.2 NEST Simulator

For users without access to biological components, the plugin supports the NEST Simulator (GNU General Public License v2) [6], a spiking neural network framework that replicates large-scale neural

Figure 2: The CL1 biocomputer platform, showing the integrated microelectrode array (MEA) and perfusion system for maintaining neuronal cultures. Adapted from [9].

dynamics in software. NEST includes templates such as Pong and provides a lower-barrier environment for prototyping. Both CL1 and NEST share standardized interfaces within our plugin: identical state/reward mappings, spike-to-game-action translation layers, and compatible data protocols. This design ensures seamless transitions between simulated and biological systems while preserving a consistent development workflow.

Figure 3: Screenshot of NEST Desktop showing a simple neural network implementation (author's own image). The interface displays a sample neuron, recorder, and stimulator.

3.2 UDP communication

Both CL1 and NEST offer extensive documentation and example Jupyter notebooks for UDP-based communication. In CL1, the CL-API registers a remote host as a "spike firehose" target, opens a socket, and listens for reply packets, which are translated into stimulation API calls[11]. NEST similarly supports UDP spike-streaming, enabling both platforms to send and receive spike events in real time.

On the Unreal Engine side, our implementation opens a UDP socket on a specified port and runs a listener thread. Incoming spike data is parsed and dispatched to the game thread via delegates or events, ensuring thread-safe updates to the game state. This bidirectional architecture allows for live stimulation, adaptive reward delivery, and dynamic environmental feedback.

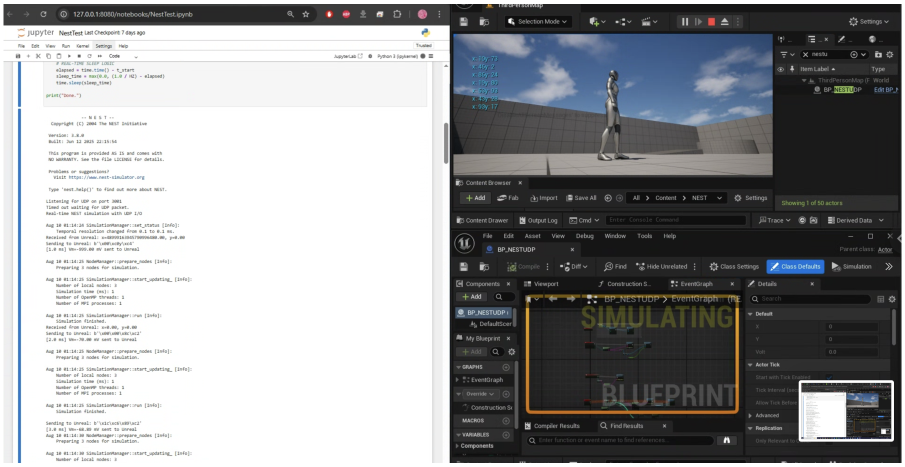

Figure 4: Example of real-time UDP communication between Unreal Engine and NEST Simulator(author's own image). Unreal Engine sends XY position data to NEST, while spike events in mV and stimulation timing pulses (1ms precision) is returned to Unreal Engine.

3.3 Plugin Functions

Beyond basic UDP send/receive, the plugin will provide higher-level functions (with Blueprint support for non-programmers), such as Send Stimulus, Get Spike Response, Send Reward Signal, Visualize Spikes, Record Session Data, Save to CSV, Stream Data to External Applications (via OSC/ Spout).

3.4 Game Templates

Inspired by OpenAI Gym[4], we include pre-built Unreal Engine templates to accelerate gamified experimentation. These are designed for reinforcement learning workflows using the Unreal Engine Learning Agents Plugin.

Example templates include:

3.4.1 3D Navigation

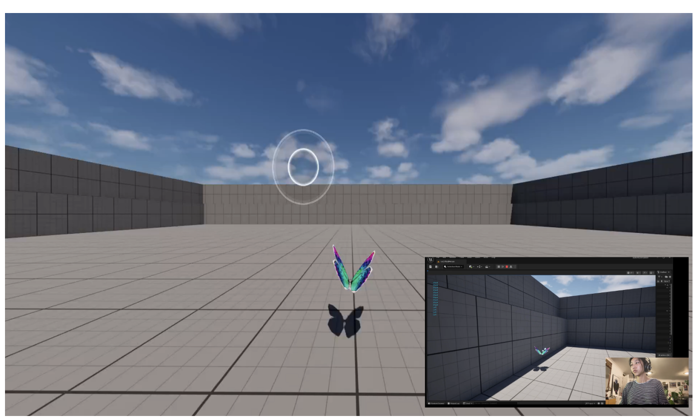

Figure 5: 3D Navigation of agent in simulated environment(author's own image). Translating spiking activities to axis action mappings

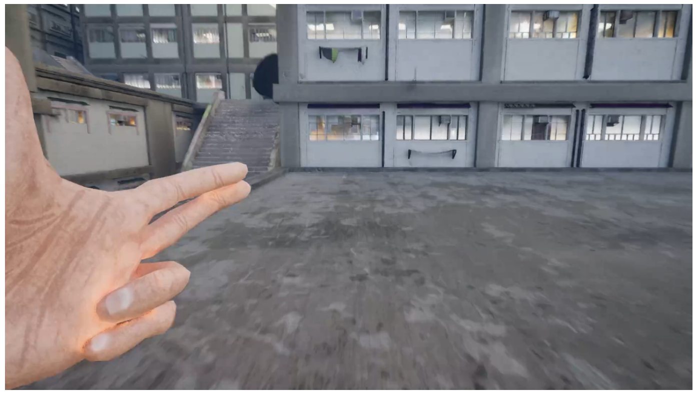

Figure 6: Adapting axis mappings and spikes to play a FPS game.

3.4.2 Team battle

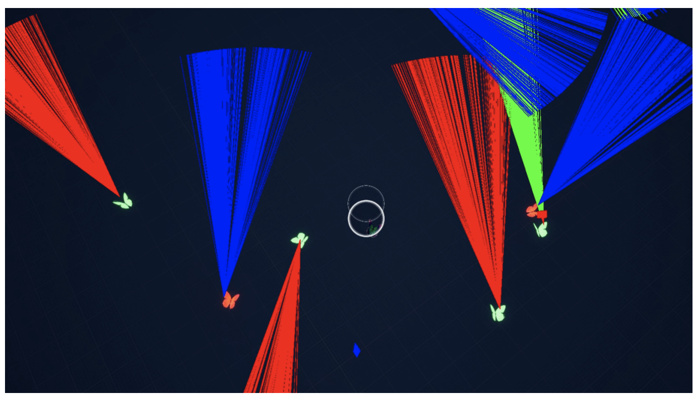

Figure 7: Team battle between two teams of agents(author's own image).

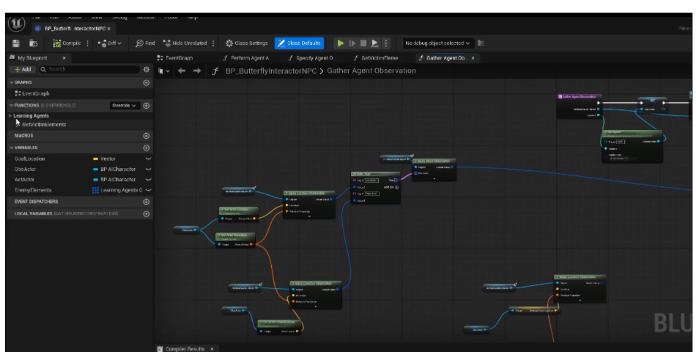

Figure 8: Set-up of reinforcement learning observations, actions, and policy using the Learning Agents Plugin.

3.4.3 Flocking / Goal-Directed Behavior ("Learning to Fly")

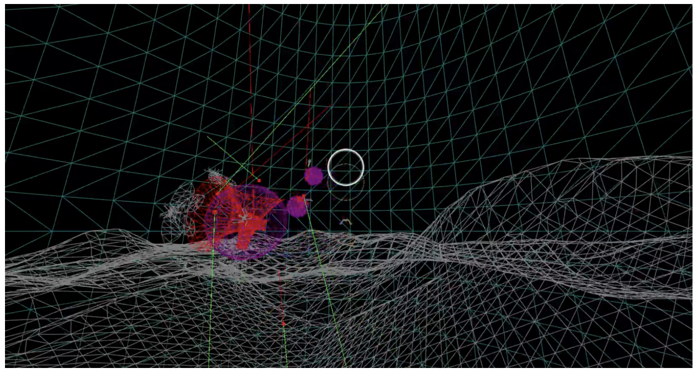

Figure 9: Goal-directed flocking/ swarming

3.4.4 Other experiments

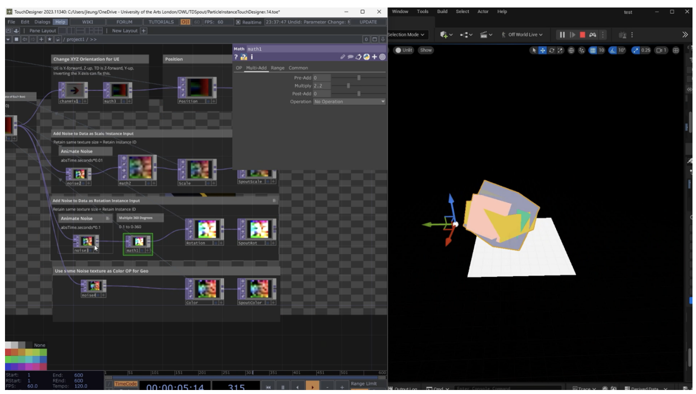

Figure 10: Supporting other data streaming or data export such as OSC or texture streaming via Spout to other software

4 Limitations

Organoid-based systems face inherent biological variability[13]— their performance in gameplay or other tasks may be inconsistent due to differences in cellular viability, maturation states, and hardware conditions [10]. Reproducibility depends critically on standardized biocomputing setups (perfusion, electrode arrays, etc.), and visible results may require extended training periods compared to digital systems.

Current brain organoid systems also face significant scaling constraints, particularly around vascularization, which limits their viable size, longevity, and functional complexity[22]. These biological limits mean that present-day organoids are thought to fall short of the neural architectures required for capacities for complex games and consciousness [8]. While this reduces the likelihood of immediate concerns, it remains essential to ensure that any future increases in neural complexity are matched with rigorous, standardized assessments for detecting relevant capacities, so that technological scaling proceeds ethically.

Another practical limitation lies in the signal-to-noise ratios of high-density microelectrode array recordings, which can constrain the reliability of closed-loop control, especially when translating noisy spike trains into discrete game actions in real time. Such variability can reduce reinforcement learning stability and complicate reproducibility across platforms.

Finally, as Nguyen warns, simplified performance metrics can drive “value capture,” where narrow reward functions eclipse broader research aims[15]. If left unchecked, this could bias both the biological system’s adaptation and the interpretation of its behavior. Mitigation strategies (reward-schema versioning, periodic metric rotation, and transparent audit logs) should remain integral to experimental design.

5 Conclusion

Assembloid Agency provides plug-and-play tools for researchers to integrate living neuronal signals into Unreal Engine environments, transforming experimental results into interactive, publishable simulations.

The open-source plugin and templates, which include reinforcement learning benchmarks (e.g. vs. DishBrain [10]) and future creative applications such as world navigation, team battles, sound synthesis, procedural environment design, LLM coupling, data streaming, enable direct comparisons between biological and artificial agents.

While this framework bridges game engines, web protocols, and biocomputing for cross-disciplinary use, its efficacy depends on addressing biological variability: organoid performance fluctuates due to culture conditions, hardware fidelity, and extended training requirements as poly-compuatational agents. We mitigate this through adaptive game design, standardized documentation, and synthetic controls (NEST simulations), ensuring robustness and playfulness at the core of the toolkit.

References

[1] Leyla Amirifar, Amir Shamloo, Rohollah Nasiri, Natan Roberto de Barros, Ze Zhong Wang, Bige Deniz Unluturk, Alberto Libanori, Oleksandr Ievglevskyi, Sibel Emir Diltemiz, Samuel Sances, Ilangko Balasingham, Stephanie K Seidlits, and Nureddin Ashammakhi. Brain-on-a-chip: Recent advances in design and techniques for microfluidic models of the brain in health and disease. *Biomaterials*, 285:121531, June 2022.

[2] Joshua Bongard and Michael Levin. Living things are not (20th century) machines: Updating mechanism metaphors in light of the modern science of machine behavior. *Front. Ecol. Evol.*, 9, March 2021.

[3] Joshua Bongard and Michael Levin. There's plenty of room right here: Biological systems as evolved, overloaded, multi-scale machines. *Biomimetics (Basel)*, 8(1):110, March 2023.

[4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI gym. *arXiv [cs.LG]*, June 2016.

[5] FinalSpark. Neuroplatform - finalspark. `https://finalspark.com/neuroplatform/`, April 2024. Accessed: 2024-9-16.

[6] Steffen Graber, Jessica Mitchell, Anno Christopher Kurth, Dennis Terhorst, Jan-Eirik Welle Skaar, Catherine Mia Schöfmann, Susanne Kunkel, Guido Trensch, Nicolai Haug, Daniel Mallett, Pishyi Yehor Andriyovich, Xavier Otazu Porter, Anthony Y Lee, and Hans Ekkehard Plesser. NEST 3.8, 2024.

[7] Edwin Hutchins. *Cognition in the Wild*. MIT Press, August 1996.

[8] Jacob Jeziorski, Reuven Brandt, John H. Evans, Wendy Campana, Michael Kalichman, Evan Thompson, Lawrence Goldstein, Christof Koch, and Alysson R. Muotri. Brain organoids, consciousness, ethics and moral status. *Seminars in Cell & Developmental Biology*, 144, Mar 2022. doi: https://doi.org/10.1016/j.semcdb.2022.03.020.

[9] Brett J Kagan. The CL1 as a platform technology to leverage biological neural system functions. *Nat. Rev. Bioeng.*, July 2025.

[10] Brett J Kagan, Andy C Kitchen, Nhi T Tran, Forough Habibollahi, Moein Khajehnejad, Bradyn J Parker, Anjali Bhat, Ben Rollo, Adeel Razi, and Karl J Friston. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23):3952–3969.e8, December 2022.

[11] Cortical Labs. Cl-api documentation (early release). `https://github.com/Cortical-Labs/cl-api-doc`, 2024. Accessed: 2025-08-09.

[12] Jenn Leung, Chloe Loewith, and Ivar Frisch. Organoid array computing: The design space of organoid intelligence. *Antikythera Digital Journal*, May 2025.

[13] Xiaofeng Liu, Zhiyuan Zhou, Yao Zhang, Hao Zhong, Xiulei Cai, and Rongfa Guan. Recent progress on the organoids: Techniques, advantages and applications. *Biomedicine & Pharmacotherapy*, 185:117942, Apr 2025. doi: https://doi.org/10.1016/j.biopha.2025.117942. URL `https://www.sciencedirect.com/science/article/pii/S0753332225001362`.

[14] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. *Playing Atari with Deep Reinforcement Learning*. DeepMind, Dec 2013. URL `https://arxiv.org/pdf/1312.5602`.

[15] Christopher Nguyen. Value capture. *J. Ethics Soc. Philos.*, 27(3), May 2024.

[16] Banjo Obayomi. Aws builder center, 2025. URL `https://builder.aws.com/content/2dbNlQiqKvUtTBV15mHBqivckmo/14-llms-fought-314-street-fighter-matches-heres-who-won`.

[17] OpenAI. Openai five, 2018. URL `https://openai.com/index/openai-five/`.

[18] Luciana Parisi. *Contagious Architecture: Computation, Aesthetics, and Space*. MIT Press, March 2013.

[19] David Silver, Thomas Hubert, Julian Schrittwieser, and Demis Hassabis. Alphazero: Shedding new light on chess, shogi, and go, Dec 2018. URL `https://deepmind.google/discover/blog/alphazero-shedding-new-light-on-chess-shogi-and-go/`.

[20] Lena Smirnova, Brian S Caffo, David H Gracias, Qi Huang, Itzy E Morales Pantoja, Bohao Tang, Donald J Zack, Cynthia A Berlinicke, J Lomax Boyd, Timothy D Harris, Erik C Johnson, Brett J Kagan, Jeffrey Kahn, Alysson R Muotri, Barton L Paulhamus, Jens C Schwamborn, Jesse Plotkin, Alexander S Szalay, Joshua T Vogelstein, Paul F Worley, and Thomas Hartung. Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish. *Front. Sci. Ser.*, 1, February 2023.

[21] G. Tesauro. Temporal difference learning and td-gammon. *ICGA Journal*, 18(2):88–88, Jun 1995. doi: https://doi.org/10.3233/icg-1995-18207. URL `https://cling.csd.uwo.ca/cs346a/extra/tdgammon.pdf`.

[22] Xingli Zhao, Zilu Xu, Lang Xiao, Tuo Shi, Haoran Xiao, Yeqin Wang, Yanzhao Li, Fangchao Xue, and Wen Zeng. Review on the vascularization of organoids and organoids-on-a-chip. *Frontiers in Bioengineering and Biotechnology*, 9, Apr 2021. doi: https://doi.org/10.3389/fbioe.2021.637048.

A Technical Appendices and Supplementary Material

Technical appendices with additional results, figures, graphs and proofs may be submitted with the paper submission before the full submission deadline (see above), or as a separate PDF in the ZIP file below before the supplementary material deadline. There is no page limit for the technical appendices.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Abstract includes all sections and methhods with clear goal to propose a UE Toolkit.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: There is ample justification of biological variability, polycomputational constraints of the subjects of study.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Paper provides a proposal for a toolkit and currently tested results. Assumptions are not fully listed as there will be variability on results based on neuronal activities and conditions of bio-computer.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Technical details such as UDP connection, RL learning agents framework have been shared and are reproducible and testable. Project is a working prototype based on publicly available code.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [NA]

   Justification: The code will be publicly shared once it is at a later stage of development. It is currently a proposal for a toolkit with already prototyped templates and protocols.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [No]

   Justification: The paper creates templates about reinforcement learning training but have no fixed and reprodicible results needed. It is for designed for both artistic and research purposes.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [NA]

   Justification: It is not an experiment or benchmark against other solutions. It is a proposal for a toolkit.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The paper provides alternatives to users without specific compute or resources.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: It recognizes the limitations of biocomputing platforms and its issues with scalability.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The paper speculates on the use of toolkit being generative for many forms of artistic and research practices, bridging multiple networking protocols, game engines and biocomputing/ simulator.

    Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: The paper uses publicly available neuronal models based on NEST simulator. The development of the toolkit does not collect user data per se.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: License included and authors credited to code.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: Documented with screenshots, process images, and prototyped templates.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: Not applicable

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: N/A

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

    Justification: Paper only used for formatting and writing, usage of LLMs is not necessary or particularly relevant to games design.s

    Guidelines:

    - The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.