

WORKSHOP ON WORLD MODELS: UNDERSTANDING, MODELLING AND SCALING

**Mengyue Yang¹, Haoxuan Li², Firas Laakom³, Xidong Feng^{4,5}, Jiaxin Shi⁴, Michael Li⁶,
Francesco Faccio^{3,7}, Jürgen Schmidhuber^{3,7}**

¹ University of Bristol, United Kingdom

² Peking University, China

³ King Abdullah University of Science and Technology, Saudi Arabia

⁴ Google DeepMind

⁵ University College London

⁶ Imperial College London

⁷ The Swiss AI Lab IDSIA, USI, SUPSI, Switzerland

1 WORKSHOP SUMMARY

The concept of the "World Models", was introduced in 1990 (Schmidhuber, 1990b; Sutton, 1991), focusing on how intelligent agents can understand and model the external world to improve decision-making and planning. World models were initially focused on modelling low-level physical quantities and interactions (Sutton, 1991; Schmidhuber, 1990a; Ha & Schmidhuber, 2018) by recurrent neural networks (RNNs) (Medsger et al., 2001). Over time, these models have evolved to capture more complex, realistic, and high-dimensional scenarios. Researchers have expanded World Models in domains like image (Ha & Schmidhuber, 2018; Hafner et al., 2020), natural language (Hao et al., 2023; Touvron et al., 2023; Achiam et al., 2023) (Large Language Model via next-token prediction), and video generation (such as SORA¹ and Genie (Bruce et al., 2024)), enabling the simulation of more interactive and sophisticated environments.

The focus of this workshop includes an exploration of the strengths and limitations of classical World Modelling backbones, such as Transformers (Robine et al., 2023; Micheli et al., 2022), RNNs, and state-space models (SSMs) (Wang et al., 2024; Hafner et al., 2020; 2023). In this workshop, our invited speakers and participants will discuss key challenges in high-dimensional spatial-temporal learning, including how these models capture long-term dependencies, handle multimodal and high-dimensional data, and balance memory with computational efficiency.

Building from these foundational topics, the workshop will also discuss the broader and evolving concept of World Models. Such as World Models in multimodal scenarios like video/text generation (Bruce et al., 2024; Touvron et al., 2023) and in more specific applications like embodied AI, healthcare and sciences. This evolution highlights the growing complexity and capabilities of World Models. By bringing together leading researchers, the workshop will cover both classical and cutting-edge techniques, and discuss how World Models can be applied across a wide range of emerging applications. Some of the fundamental questions that this workshop aims to address are:

1.1 UNDERSTANDING THE WORLD AND EXTRACTING KNOWLEDGE

- What constitutes meaningful information from observations, and how can we enhance causal understanding and reasoning abilities in models?
- What are the theoretical foundations that guide the construction of World Models?
- How do ethical considerations factor into the development and deployment of World Models, especially in critical applications?

¹<https://openai.com/index/sora/>

1.2 WORLD MODEL TRAINING AND EVALUATION

- What are the strengths, limitations, and challenges of current modelling architectures such as Transformers, RNNs, and SSMs?
- What are the strengths, limitations, and challenges of current training algorithms such as autoregressive training (Sutskever et al., 2014), diffusion modelling (Sohl-Dickstein et al., 2015), and normalizing flow (Papamakarios et al., 2021)?
- What are the challenges and best practices for constructing and curating datasets that effectively train and evaluate World Models?

1.3 SCALING WORLD MODELS ACROSS LANGUAGE, VISION, AND CONTROL

- How can a pre-trained large language model serve as the World Model in language?
- What are the limitations and challenges of current vision generation tasks, such as image generation and video generation?
- What are the benefits and challenges of integrating World Models with existing reinforcement learning and control frameworks (e.g. Model-based RL (Sutton, 1991))?
- How can we build a unified Multimodal World Model across domains, such as text-image, or text-image-action models?
- What challenges arise when scaling World Models?

1.4 WORLD MODELS IN GENERAL DOMAINS: EMBODIED AI, HEALTHCARE, NATURAL AND SOCIAL SCIENCES, AND BEYOND

- How to construct a World Model in general domains?
- How can applications in general domains benefit from World Models?

Importantly, our workshop covers the widest range of topics related to World Models, including understanding, and modelling, as well as closely aligning with cutting-edge generative AI and broader applications such as embodied AI. We have invited researchers from both academia and industry with impressive expertise in the various related fields, including the most advanced areas of generative AI. Meanwhile, our workshop widely targets AI researchers, industry professionals, and students interested in World Models, generative AI, reinforcement learning and related applications.

Participants should have a basic understanding of generative models and reinforcement learning concepts. Familiarity with recent advancements in both fields will be beneficial but not mandatory. We also welcome submissions from researchers in the natural sciences (e.g., physics, chemistry, biology) and social sciences (e.g., pedagogy, sociology) to offer attendees a more comprehensive perspective. In summary, our topics of interest mainly include, but are not limited to:

- Understanding World Rules: Exploring how World Models capture environment dynamics, causal understanding, spatial-temporal patterns, and theoretical foundations for simulation and prediction.
- World model training and evaluation: strengths, limitations, and challenges of current modelling architectures (e.g. Transformers, RNNs, and SSMs), training algorithms (autoregressive training, diffusion modelling, and normalizing flow) and dataset construction.
- Scaling World Models across language, vision, and control: Investigating how integrating visual, auditory, and textual data improves realism World Models.
- World Models in general domains: Exploring World Models in robotics, AI, healthcare, natural and social sciences, and beyond to improve prediction and decision-making.

2 WORKSHOP SCHEDULE

Our one-day workshop mainly includes 3 poster sessions, 8 invited talks, 2 oral sessions with 6 contributed talks in total, and a panel discussion. We are glad to announce that **10 confirmed top-tier researchers as speakers and panelists, 9 of them have confirmed to attend in person**. In addition, we plan to organize a lunch buddy, aiming to encourage a more interactive and inclusive environment connecting early-career researchers with senior scholars (see below for more details). For the contributed paper sessions, regarding the recent surge in publications in related areas and the success of similar workshops, we project over 250 paper submissions and over 1,500 participants.

2.1 TENTATIVE WORKSHOP SCHEDULE

08:15 am - 09:20 am: Poster Session I

09:20 am - 09:30 am: Opening Remarks

09:30 am - 10:00 am: Invited talk 1

10:00 am - 10:30 am: Invited talk 2

10:30 am - 11:00 am: Invited talk 3

11:00 am - 11:30 am: Oral Presentation I (Total 3 talks, 10 mins for each)

11:30 am - 01:00 pm: Lunch Break + Lunch Buddy (see below for more details) + Poster Session II

01:00 pm - 01:30 pm: Invited talk 4

01:30 pm - 02:00 pm: Invited talk 5

02:00 pm - 02:30 pm: Invited talk 6

02:30 pm - 03:00 pm: Oral Presentation II (Total 3 talks, 10 mins for each)

03:00 pm - 03:30 pm: Coffee Break + Poster Session III

03:30 pm - 04:00 pm: Panel discussion

04:00 pm - 04:30 pm: Invited talk 7

04:30 pm - 05:30 pm: Invited talk 8

05:30 pm - 06:00 pm: Social

Lunch Buddy Plan: Building on an idea introduced at the NeurIPS 2022 NeurReps workshop, we plan to launch a "Lunch Buddy" program, which aims to encourage a more interactive and inclusive environment. This initiative is designed to connect early-career researchers with senior scholars, as well as individuals from diverse communities, over lunch. Participants will have the opportunity to sign up for their interest online before workshop day and provide information about their research interests and levels of experience. During the morning sessions of the workshop, participants will confirm their attendance. Participants will then be randomly paired, and we will offer a list of suggested discussion topics to guide their lunchtime conversations.

2.2 CONFIRMED INVITED SPEAKERS AND PANELIST

Jeff Clune (Confirmed as speaker and panelist). Professor at University of British Columbia; Canada CIFAR AI Chair; Faculty Member at Vector Institute; Senior Research Advisor at DeepMind. His research focuses on deep-learning and reinforcement learning, especially in creating open-ended algorithms and AI-generating algorithms. He has extensive presentation experience in top-tier conferences such as CORL 2021, ICML 2020, and NeurIPS 2019. Personal Website: <http://jeffclune.com/>

Chelsea Finn (Confirmed as speaker and panelist). Assistant Professor in Computer Science and Electrical Engineering at Stanford University and co-founder of Pi. Her research interests lie in the capability of robots and other agents to develop broadly intelligent behavior through

learning and interaction. To this end, her work has pioneered end-to-end deep learning methods for vision-based robotic manipulation, meta-learning algorithms for few-shot learning, and approaches for scaling robot learning to broad datasets. He has extensive presentation experience in top-tier conferences such as ICML, ECCV, CVPR, ICRA, CoRL and NeurIPS. Personal Website: <https://ai.stanford.edu/~cbfinn/>

Tim Rocktäschel (Confirmed as speaker and panellist). Professor in the Department of Computer Science at University College London (UCL); Senior Staff Research Scientist and the Open-Endedness Team Lead at Google DeepMind; Fellow of the European Laboratory for Learning and Intelligent Systems (ELLIS). His research spans on Artificial General Intelligence, Open-Endedness, and Self-Improvement. He has received two Best Paper Awards at ICML 2024. Personal Website: <https://rockt.github.io/>

Stefano Ermon (Confirmed as speaker and panellist). Associate Professor in the Department of Computer Science at Stanford University. His research is in machine learning and generative AI, especially for developing principled methods motivated by concrete real-world applications and problems of broad societal relevance. He has received the Best Paper Awards (Outstanding Paper Awards) in ICML 2024, ICLR 2022, ICLR 2021 and AAAI 2017. Personal Website: <https://cs.stanford.edu/~ermon/>

Jakob Foerster (Confirmed as speaker). Associate Professor in the Department of Engineering Science at University of Oxford; Research Scientist in Meta AI. His research interests span on Deep and Multi-agent Reinforcement Learning. Personal Website: <https://www.jakobfoerster.com/>

Furong Huang (Confirmed as speaker). Associate Professor in the Department of Computer Science at the University of Maryland. Her research interest is in trustworthy machine learning, AI for sequential decision-making, and high-dimensional statistics. She has extensive experience in delivering conference presentations, including oral talks at major events and workshops such as MLSys 2023, ICML 2023, and ICLR 2023, CVPR 2024. Personal Website: <https://furong-huang.com/>

Xiaolong Wang (Confirmed as speaker). Assistant Professor at UC San Diego in the ECE department. His research focuses on computer vision, reinforcement learning and robotics, especially for learning 3D and dynamics representations through videos and physical robotic interaction data. He has extensive experience in delivering presentations in workshops for top conferences such as CVPR, CoRL, ICLR, ECCV and ICRA. Personal Website: <https://xiaolonw.github.io/>

Tom Everitt (Confirmed as speaker). Staff Research Scientist in Google DeepMind. His research focuses on the causality and agent problem for AGI Safety, world understanding, and modelling. He has received the Honorable Mention Outstanding Paper Award in ICLR 2024. He has extensive experience in delivering conference presentations, including oral talks and tutorials at major events such as UAI 2023, AAAI 2023, and ICLR 2024. Personal Website: <https://www.tomeveritt.se/>

Kun Zhang (Confirmed as panellist). Associate professor at Carnegie Mellon University, and professor at Mohamed bin Zayed University of Artificial Intelligence (MBZUAI). His research interests lie in causal discovery, causal representation learning, and causality-based learning, aiming to make hidden entities and causal processes transparent for the purpose of automated scientific discovery, and optimal decision-making. Personal Website: <https://www.andrew.cmu.edu/user/kunz1/index.html>

David Ha (Confirmed as panellist). Research Scientist at Sakana AI. He used to be the Head of Strategy at Stability AI, and one of the top minds working in AI today. He previously worked as a research scientist in the Brain team at Google. David is particularly interested in evolution and complex systems, and his research explores how intelligence may emerge from limited resource constraints. He is now working on world models and interested in studying how intelligence might have emerge from limited resource constraints. Personal Website: <https://otoro.net/ml/>

3 ORGANIZING TEAM

Our team includes researchers specializing in World Models, representation learning theory, causality, RL, and robotics, all of which are closely aligned with the workshop’s theme. **Importantly, our team is fortunate to have founders of World Models: Jürgen Schmidhuber as our workshop co-organizer and David Ha as invited keynote speaker.** Our team not only brings a wealth of experience across various research domains but also showcases remarkable diversity in backgrounds, and genders. With 8 members from 7 different institutions across 5 countries, including North America, Europe, the Middle East, and Asia, we are well-positioned to integrate a broad range of perspectives and expertise into the workshop. Every team member is fully committed to ensuring a rewarding experience for contributors, speakers, and attendees alike.

Our team members have extensive experience in organizing workshops and symposiums, making us well-prepared to handle workshops organization at large conferences. The organized important events of our team are listed below.

- Universal Learning Algorithms and Optimal Search Workshop at NeurIPS 2022
- Causality and Large Model Workshop at NeurIPS 2024;
- Causal Representation Learning Workshop at NeurIPS 2024;
- Artificial Intelligence with Causal Techniques Workshop at AAAI 2025;
- Causal Structure Learning from Event Sequences and Prior Knowledge Competition at NeurIPS 2023;
- Causal Representation Learning Workshop at ICDM 2024;
- Causal Recommendation: Progresses and Future Directions Tutorial at SIGIR 2023;
- Causality and Large Models Tutorial at ACML 2024;
- Theoretical Foundations and Applications of Deep Generative Models Workshop Keynote Speaker at ICML 2018;
- Gaze Meets ML Workshop Keynote Speaker at NeurIPS 2022;
- Metacognition in AI Workshop Keynote Speaker at NeurIPS 2021.

3.1 ORGANIZER INTRODUCTIONS

Mengyue Yang is a Lecturer (equals to US Assistant Professor) in AI at University of Bristol. She obtained her Ph.D. from University College London, under the supervision of Professor Jun Wang. Her research interests are causality, multi-agent systems and reinforcement learning. She has published several first-author papers in top-tier AI conferences and journals such as NeurIPS, CVPR, KDD, SIGIR, WWW, and ACM TOIS. She has been recognized as a Rising Star in AI by KAUST. She has extensive experience in organizing workshops in conferences. She was co-organizer of the Causal Representation Learning Workshop at ICDM 2024, Causal Structure Learning from Event Sequences and Prior Knowledge Competition at NeurIPS 2023, Causality and Large Model Workshop at NeurIPS 2024 and Artificial Intelligence with Causal Techniques Workshop at AAAI 2025.

Haoxuan Li is a Ph.D. candidate in Center for Data Science, Peking University. His research interests span from causal machine learning, decision-making theory, trustworthy AI, and large language models. He has published more than 40 papers in top-tier AI conferences such as ICML, NeurIPS, ICLR, AAAI, IJCAI, SIGKDD, WWW, SIGIR, etc., and received the NSFC Young Scientists Fund (2024). Moreover, he has organized the AAAI 2025 Workshop on Artificial Intelligence with Causal Techniques, and co-organized the NeurIPS 2024 Workshop on Causal Representation Learning.

Firas Laakom is a postdoctoral researcher at GenAI in KAUST, mentored by Professor Jürgen Schmidhuber. Prior to this, he obtained his Ph.D. from Tampere University, Finland under the supervision of Professor Moncef Gabbouj. He has co-authored several first-author papers at top-tier AI conferences and journals. His research interests include General AI, neural networks, learning theory, and sequential modeling.

Xidong Feng is a research scientist at Google DeepMind. His research spans over Large Language Model, Reinforcement Learning and Multi-agent Learning. He has published over 10 papers in top

AI conferences or journals like NeurIPS, ICML and JMLR. Previously, he obtained his Ph.D. at Computer Science, University College London, advised by Prof. Jun Wang.

Jiaxin Shi is a research scientist at Google DeepMind. Previously, he was a postdoctoral researcher at Stanford and Microsoft Research New England. He obtained Ph.D. from Tsinghua University under the advisement of Professor Jun Zhu. His research interests broadly involve probabilistic and algorithmic models for learning as well as the interface between them. Jiaxin served as an area chair for NeurIPS and AISTATS. He is the workflow chair of AISTATS 2024. He is a recipient of Microsoft Research PhD fellowship. His first-author paper was recognized by a NeurIPS 2022 outstanding paper award.

Michael Li is a Department Lecturer/Chapman Fellow at Department of Mathematics at Imperial College London. He was previously a Senior Research Fellow at the Gatsby Computational Neuroscience Unit. He earned his Ph.D. from the University of Oxford under the supervision of Professor Dino Sejdinovic. His research focuses on understanding the computational and statistical mechanisms underlying multi-stage learning algorithms, which form the foundational training protocols for many AI and machine learning models, including large language models, diffusion, and causal inference. His work has been recognized with an ICML 2019 Best Paper Honorable Mention and a NeurIPS 2022 Oral Presentation Award. He has also served as a reviewer for leading conferences such as NeurIPS, ICML, and ICLR, and has actively organized workshops on kernel learning theory and tensor learning in University College London Gatsby Unit and Oxford University.

Francesco Faccio is a postdoctoral researcher at IDSIA, mentored by Prof. Jürgen Schmidhuber. Before this, he completed a PhD in the same lab. Before his PhD, he earned bachelor's and master's degrees in Mathematical Engineering at Politecnico di Milano. His research interests lie at the intersection of General AI, Reinforcement Learning, and AI for Science, focusing on developing artificial scientists to automate scientific research using artificial curiosity and meta-learning techniques. He has authored 11 publications in top-tier AI conferences and journals, with four receiving oral presentations at NeurIPS, ICML, and AAI, and one winning the Best Paper Award at the NeurIPS Ro-FoMo Workshop. In the past two years, he was a visiting researcher at KAUST, where he developed and authored projects on General AI, AI for Chemistry, Trustworthy Large Language Models, Car Racing Optimization, and Coral Restoration. He also co-organized the KAUST Rising Stars in AI Symposium in 2023 and 2024.

Jürgen Schmidhuber is the Scientific Director of Swiss AI Lab, IDSIA and the co-chair of the Center of Excellence in Generative AI at KAUST. The main goal of Professor Jürgen Schmidhuber has been to build a self-improving Artificial Intelligence (AI) smarter than himself. His lab's Deep Learning Neural Networks (NNs) based on ideas published in the "Annus Mirabilis" 1990-1991 have revolutionised machine learning and AI. He was one of the first to work on LSTM, feedforward NNs on GPUs, DanNet, deep NN for medical imaging, GANs, Transformers. His research group also established the fields of mathematically rigorous universal AI and recursive self-improvement in meta-learning machines that learn to learn (since 1987). He also generalized algorithmic information theory and the many-worlds theory of physics. He is recipient of numerous awards, author of about 400 peer-reviewed papers, and Chief Scientist of the company NNAISENSE, which aims at building the first practical general purpose AI.

3.2 PROGRAM COMMITTEE

By leveraging the collaboration between seven institutions, we have established a strong network of potential reviewers, ensuring that *each submission will receive at least three high-quality reviews* and *each reviewer will not be allocated by more than three papers*. We are pleased to confirm that **we have already secured approximately three-fourths of the required reviewers** through an internal sign-up process at three of the participating institutions. While this is a promising start, we are continuing to seek additional qualified reviewers to guarantee a fair and thorough review process. Several members of our organizing team will act as area chairs, utilizing their expertise in specific causal subfields to manage the review process efficiently.

4 OTHER STATEMENTS

4.1 SUBMISSION GUIDELINES

Submissions should follow the main track ICLR format for 4-9 pages, shorter working-in-progress submissions with proof-of-concept demonstrations with live demos, code, and blog posts are welcome. We encourage the submission of LaTeX templates. Detailed submission guidelines and templates will be provided on the workshop website.

All the accepted papers will be announced by 5th March 2025 and the accepted papers will be made available on the workshop website. The workshop will not have any proceedings. Meanwhile, to incentivize participants to submit new and unique work, our workshops will not accept the work that has already been published at other machine learning conferences (including the main ICLR conference).

We will have three poster presentation sessions during the workshop, giving authors and attendees plenty of time to discuss ideas. These sessions will help authors improve their work and give attendees a chance to see new perspectives. We will also award a Best Paper to the most innovative or unique contribution at the workshop.

4.2 ONLINE RESOURCES

Due to the rapid growth in the scale of the conference and the number of participants, it is difficult for us to accurately estimate whether the venue can accommodate all attendees and support sufficient interaction. Therefore, we will add more online activities and resources to ensure broader participation.

Recording of Talks: We will ensure that all workshop talks are recorded and made available to those unable to attend in person. This will not only benefit participants in different time zones but also provide resources for further study and discussion after the event.

Online Resources: We plan to post the presentation title and abstract online on our website before the workshop days. Moreover, short working papers, posters, industry demos and relevant discussion materials will be posted online after accepting. These resources will be accessible through a dedicated website or platform, ensuring a broader audience can engage with the presented work.

Maintaining a Content-Rich Website: We will create and maintain a dedicated website or platform that will continuously update the workshop's content, including videos, papers, and discussion boards, ensuring that both participants and a wider audience have easy access to the resources.

Online Participation: All of workshop presentations and panel discussions will be held in person; however, in the event of unforeseen circumstances, we will offer contingency plans to provide online support for talks.

Online Interactive Q&A and Forum: We will provide an online forum or Q&A session where those who missed the workshop can interact with the speakers, submit questions, and provide feedback, enhancing the overall level of interaction.

4.3 FUNDING AND SPONSOR

Funding upon acceptance, we will solicit sponsorships from companies such as Apple, Meta and Google to improve the accessibility and inclusivity of our event. In particular, we will cover the registration fees of participants who is prioritizing members from under-represented groups.

4.4 DIVERSITY STATEMENT

During the selection of organisers and speakers, we actively encouraged all forms of diversity. The expertise of invited speakers ranges from theory to applications, including World Models, RL to generative models and robotics, while also achieving gender parity and covering different research institutions, nationalities, scientific backgrounds and levels of seniority. The full scale of scientific seniority is covered in both organisation committees and speakers, including Ph.D. candidate, postdoc researchers, assistant professors, full professors as well as researchers in industry.

4.5 RELEVANCE AND IMPACT

The past event included the NeurIPS 2024 Workshop on Open-World Agents: Synergizing Reasoning and Decision-Making in Open and Interactive Environments. IROS 2023 Workshop on World Models and Predictive Coding in Cognitive Robotics. Previous workshops have focused on specific areas, such as robotics or the agent problem. In contrast, our workshop aims to present a more comprehensive discussion on the technology of world models. We will include classical world model methods—like sequence modelling, and initiate broader discussions on the evolution of world models, such as multimodal modelling, generative AI and their integration with more general application areas. Our goal is to connect AI researchers widely and encourage a holistic exchange.

The topic broadly includes sequence modelling, causality, reinforcement learning, and planning, and is also closely aligned with cutting-edge multimodal modelling and generative model, which is expected to attract significant attention from researchers. At the same time, our theme encourages participation from general domains, such as world models in natural sciences (chemistry, biology) and scientific discovery. This broad discussion will foster cross-disciplinary collaboration and provide a platform for innovative AI applications across various fields.

SELECTED REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative interactive environments. In *Forty-first International Conference on Machine Learning*, 2024.
- David Ha and Jürgen Schmidhuber. World models. *Advances in Neural Information Processing Systems*, 2018.
- Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. Reasoning with language model is planning with world model. *arXiv preprint arXiv:2305.14992*, 2023.
- Larry R Medsker, Lakhmi Jain, et al. Recurrent neural networks. *Design and Applications*, 5(64-67): 2, 2001.
- Vincent Micheli, Eloi Alonso, and François Fleuret. Transformers are sample-efficient world models. *arXiv preprint arXiv:2209.00588*, 2022.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- Jan Robine, Marc Höftmann, Tobias Uelwer, and Stefan Harmeling. Transformer-based world models are happy with 100k interactions. *arXiv preprint arXiv:2303.07109*, 2023.
- Jürgen Schmidhuber. An on-line algorithm for dynamic reinforcement learning and planning in reactive environments. In *1990 IJCNN international joint conference on neural networks*, pp. 253–258. IEEE, 1990a.
- Jürgen Schmidhuber. Making the world differentiable: On using self-supervised fully recurrent neural networks for dynamic reinforcement learning and planning in non-stationary environments, 1990b.

- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. PMLR, 2015.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *arXiv preprint arXiv:1409.3215*, 2014.
- Richard S Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Xiao Wang, Shiao Wang, Yuhe Ding, Yuehang Li, Wentao Wu, Yao Rong, Weizhe Kong, Ju Huang, Shihao Li, Haoxiang Yang, et al. State space model for new-generation network alternative to transformers: A survey. *arXiv preprint arXiv:2404.09516*, 2024.