

Danijar Hafner

mail@danijar.com

[Google Scholar](#) [Website](#) [Github](#) [Twitter](#)

I aim to develop autonomous AI that makes decisions given only sporadic human feedback or instruction. To accomplish this, I research scalable world models that equip AI with a rich perception of the world, the ability to imagine future outcomes, to explore autonomously, and to achieve distant goals through hierarchical planning. I have been fortunate to contribute to this challenge through 21 publications at top venues. Our world models taught a robot dog to walk in just 1 hour from scratch in the real world, leading to interviews with the Daily Mail, MIT Tech Review, Tech Crunch, and other magazines.

EDUCATION	University of Toronto PhD Computer Science Advisor: <i>Jimmy Ba</i>	2018–2023
	University of California, Berkeley Visiting PhD Student Advisor: <i>Pieter Abbeel</i>	2022–2023
	University College London MRes Computational Statistics and Machine Learning Advisors: <i>Timothy Lillicrap and Karl Friston</i>	2017–2018
	Hasso Plattner Institute BSc IT Systems Engineering	2013–2017
INDUSTRY	DeepMind Research Scientist Intern	2022–2023
	Google Research Brain Team 6 years of industry research experience in Mountain View, London, Toronto, and San Francisco. Fun Fact: <i>I completed 12 Google internships in total.</i>	2016–2022
FUNDING	Vanier Scholarship Canada's most prestigious doctoral scholarship (\$150,000). Top 13 across all sciences out of 172 awardees and many applicants.	2020–2023

INTERVIEWS	The Future of Development Podcast	upcoming
	TalkRL: The Reinforcement Learning Podcast	upcoming
	Daily Mail Link	07/2022
	MIT Technology Review Link	07/2022
	Tech Crunch Link	07/2022
	Tech Xplore Link	07/2022
	Towards Data Science Podcast Link	01/2022
	TalkRL: The Reinforcement Learning Podcast Link	05/2020
INVITED PANELIST	ICLR Open-Ended Agent Learning Workshop	04/2022
	Data ICMC Open-Endedness Panel	10/2021
	ICML Unsupervised RL Workshop	07/2021
	ICLR Never Ending RL Workshop	05/2021
	ICLR Self-Supervised RL Workshop	05/2021
SELECTED INVITED TALKS	I regularly receive talk invitations but decline most of them to leave more time for research.	
	ICLR Open-Ended Agent Learning Workshop	04/2022
	ICAPS AI Planning & RL Workshop Opening Talk	08/2021
	ICML Unsupervised RL Workshop	07/2021
	ICLR Never Ending RL Workshop	05/2021
	ICLR Self-Supervised RL Workshop	05/2021
	DeepMind Embodied Intelligence Group	10/2020
	DeepMind Montreal Team	10/2020
	Intel AI Labs	07/2020
Facebook AI Research	11/2019	

Wellcome Trust Centre	FIL	07/2019
DeepMind	Meta Learning Group	06/2019
Berkeley AI Research	BAIR	06/2019
Wellcome Trust Centre	FIL	06/2018
IBM Research		02/2018

MEDIA
COVERAGE

DayDreamer: World Models for Physical Robot Learning
Daily Mail, MIT Technology Review, TechCrunch, Tech Xplore, Synced, SingularityHub, ZME Science, Technology Org, India Times, Analytics India Mag, MarkTechPost, NEWS7g, Actula, I Programmer

Deep Hierarchical Planning from Pixels
MarkTechPost

Planning to Explore via Self-Supervised World Models
VentureBeat, Synced, Yannic Kilcher

Dream to Control: Learning Behaviors by Latent Imagination
VentureBeat, State of AI Report 2020, Yannic Kilcher, Towards Data Science, Siecle Digital, India AI, MC.AI, Towards AI,

Learning Latent Dynamics for Planning from Pixels
Two Minute Papers, VentureBeat, Packt, Neurohive, Towards Data Science, Edgy

JOURNAL &
CONFERENCE
PAPERS

See [Google Scholar](#) for a complete list of publications.

Deep Hierarchical Planning from Pixels NeurIPS 2022
D Hafner, KH Lee, I Fischer, P Abbeel

Learning Robust Dynamics through Variational Sparse Gating NeurIPS 2022
AK Jain, SK Sujit, S Joshi, V Michalski, D Hafner, SE Kahou

World Models for Physical Robot Learning CoRL 2022
P Wu, A Escontrela, D Hafner, K Goldberg, P Abbeel

- Masked World Models for Visual Control** CoRL 2022
Y Seo, D Hafner, H Liu, F Liu, S James, K Lee, P Abbeel
- Benchmarking the Spectrum of Agent Capabilities** ICLR 2022
(DRLW oral)
D Hafner
- Discovering and Achieving Goals via World Models** NeurIPS 2021
(URLW oral & SSLW oral)
R Mendonca, O Rybkin, K Daniilidis, D Hafner, D Pathak
- Clockwork Variational Autoencoders** NeurIPS 2021
V Saxena, J Ba, D Hafner
- Information is Power: Intrinsic Control via Information Capture** NeurIPS 2021
N Rhinehart, J Wang, G Berseth, J Co-Reyes, D Hafner, C Finn, S Levine
- Latent Skill Planning for Exploration and Transfer** ICLR 2021
K Xie, H Bharadhwaj, D Hafner, A Garg, F Shkurti
- Mastering Atari with Discrete World Models** ICLR 2021
D Hafner, T Lillicrap, M Norouzi, J Ba
- Sophisticated Inference** Neural Communication 2021
D Tran, M Dusenberry, M van der Wilk, D Hafner
- Planning to Explore via Self-Supervised World Models** ICML 2020
R Sekar, O Rybkin, K Daniilidis, P Abbeel, D Hafner, D Pathak
- Dream to Control: Learning Behaviors by Latent Imagination** (oral) ICLR 2020
D Hafner, T Lillicrap, J Ba, M Norouzi
- A Deep Learning Framework for Neuroscience** Nature Neuroscience 2019
Richards, Lillicrap, Beaudoi, Bengio, Bogacz, Christensen, Clopath, Costa, Berker, Ganguli, Gillon, Hafner, Kepecs, Kriegeskorte, Latham, Lindsay, Miller, Naud, Pack, Poirazi, Roelfsema, Sacramento, Saxe, Scellier, Schapiro, Senn, Wayne, Yamins, Zenke, Zylberberg, Therien, Kording

	Bayesian Layers: A Module for Neural Network Uncertainty	NeurIPS 2019
	D Tran, M Dusenberry, M van der Wilk, D Hafner	
	Noise Contrastive Priors for Functional Uncertainty	UAI 2019
	D Hafner, D Tran, A Irpan, T Lillicrap, J Davidson	
	Learning Latent Dynamics for Planning from Pixels	ICML 2019
	D Hafner, T Lillicrap, I Fischer, R Villegas, D Ha, H Lee, J Davidson	
	Routing for On-Street Parking Search using Probabilistic Data	AI Communications
	T Friedrich, MS Krejca, R Rothenberger, T Arndt, D Hafner, T Kellermeier, S Krogmann, A Razmjou	
	Sample-Efficient Reinforcement Learning with Stochastic Ensemble Value Expansion (long oral)	NeurIPS 2018
	J Buckman, D Hafner, G Tucker, E Brevdo, H Lee	
	Sim-to-Real: Learning Agile Locomotion For Quadruped Robots	RSS 2018
	J Tan, T Zhang, E Coumans, A Iscen, Y Bai, D Hafner, S Bohez, V Vanhoucke	
	Probabilistic Routing for On-Street Parking Search	ESA 2016
	T Arndt, D Hafner, T Kellermeier, S Krogmann, A Razmjou, MS Krejca, R Rothenberger, T Friedrich	
WORKSHOP PAPERS	Evaluating Long-Term Memory in 3D Mazes	DRLW 2022
	J Pasukonis, T Lillicrap, D Hafner	
	Guiding Exploration Towards Impactful Actions	DRLW 2022
	V Saxena, J Ba, D Hafner	
	Learning Robust Dynamics through Variational Sparse Gating	DRLW 2021
	AK Jain, SK Sujit, S Joshi, V Michalski, D Hafner, SE Kahou	
	Action and Perception as Divergence Minimization	DRLW 2020
	D Hafner, P Ortega, J Ba, T Parr, K Friston, N Heess	

Danijar Hafner

	Evaluating Agents Without Rewards (oral)	BARL 2020
	B Matusch, J Ba, D Hafner	
	Modulated Policy Hierarchies	DRLW 2018
	A Pashevich, D Hafner, J Davidson, R Sukthankar, C Schmid	
	Learning Hierarchical Information Flow with Recurrent Neural Modules	NeurIPS 2017
	D Hafner, A Irpan, J Davidson, N Heess	
OUTREACH & LEADERSHIP	State, Action, and Time Abstraction in Reinforcement Learning	03/2023
	Workshop Organizer Co-Lead, Bellairs Institute, Barbados Notable Attendees: <i>Yann LeCun, David Silver</i>	
	Task-Agnostic Reinforcement Learning	05/2019
	Workshop Organizer Lead, ICLR Facilitated research contributions by over 200 scientists.	
	Deep Learning at Supercomputer Scale	12/2017
	Workshop Organizer, NeurIPS	
	Robot Learning Lab Tours	07/2022
	Introduced middle school and high school students to robotics by showing them the lab, live robot demos, and learning materials.	
	High School Student Research Mentoring	2019–2021
	Mentored two high school students with regular meetings on reinforcement learning research, leading to a workshop talk.	
	Berkeley Deep Reinforcement Learning Bootcamp	09/2017
	Helped facilitate this learning event as a volunteer.	
TEACHING	Interactive and Embodied Learning CS422	03/2022
	Guest lecture, Stanford University. Nick Haber and Fei-Fei Li's course.	
	AI Automated Planning CSC2542	03/2021
	Guest lecture, University of Toronto. Sheila McIlraith's course.	

TensorFlow for Deep Learning Research CS20 Spring 2018

Course Advisor, Stanford University.

Co-designed course structure and materials.

PhD Intern Research Conference Workshop Google 07/2017

Interactive tutorial on recurrent neural networks for 380 PhD students attending in person.

TensorFlow for Deep Learning Research CS20SI 03/2017

Guest lecture, Stanford University. Chip Huyen's course.

High School Programming Instructor 2015

Fun fact: *As a high school student, my teacher let me take over the class for 3 months to teach my classmates C++.*

ADVISING

Independents

Jurgis Pasukonis (next: Research Engineer at DeepMind)

PhDs

Wilson Yan (UC Berkeley)

Yonggyo Seo (KAIST)

Oleh Rybkin (U Penn)

Russell Mendonca (CMU)

Arnav Kumar Jain (MILA)

Kevin Xie (U Toronto)

Homanga Bharadhwaj (next: PhD at CMU)

Ramanan Sekar (next: ML Engineer at Qualcomm Research)

Masters

Vaibhav Saxena (next: PhD at Georgia Tech)

Minghan Li (next: PhD at U Waterloo)

Harini Kannan (next: Stealth Startup)

Manav Choudhary (next: Applied Science Lead at Zalando)

Undergraduates

Calvin Luo (next: Google AI Resident)

Sam Sinha (next: AI Resident at FAIR)

Highschoolers

Kevin Frans (next: OpenAI intern)

Brendon Matusch (next: Stanford BSc)

SOFTWARE

Top 1% overall on [Stackoverflow](#).

Open source projects with 5000 stars and 600 forks on [Github](#).

Authored 150,000 lines of Python and 25,000 lines of C++ code.

Crafter [Link](#) 200 Stars

Benchmarking the Spectrum of Agent Capabilities

DreamerV2 [Link](#) 600 Stars

Mastering Atari with Discrete World Models

Dreamer [Link](#) 500 Stars

Dream to Control: Learning Behaviors by Latent Imagination

PlaNet [Link](#) 1100 Stars

Learning Latent Dynamics for Planning from Pixels

Python Handout [Link](#) 2000 Stars

Turn Python scripts into handouts with Markdown and figures

BatchPPO [Link](#) 900 Stars

Efficient Batched Reinforcement Learning in TensorFlow

PATENTS

Training Reinforcement Learning Agents to Learn 11/2020

Farsighted Behaviors by Predicting in Latent Space

D Hafner, M Norouzi, T Lillicrap

System and Methods for Pixel Based Model Predictive 05/2019

Control

D Hafner

Sample-Efficient Reinforcement Learning 05/2019

D Hafner, J Buckman, H Lee, E Brevdo, G Tucker

Batched Reinforcement Learning 08/2018

D Hafner

Danijar Hafner

REVIEWING	Neural Information Processing Systems (NeurIPS)	2019, 2020
	International Conference on Machine Learning (ICML)	2019, 2020
	International Conference on Learning Representations (ICLR)	2019
	International Journal of Computer Vision (IJCV)	2019
	Biological Cybernetics Journal (BICY)	2018
	NeurIPS Deep Reinforcement Learning (DRLW)	2018, 2019, 2020, 2021, 2022
	ICLR Task-Agnostic Reinforcement Learning (TARL)	2019
	ICML Multi-Task and Lifelong Reinforcement Learning (MTLRL)	2019
COMPETITIONS	HPI Data Analytics and ML Hackathon 1st Place	06/2016
	Used density estimation over document embeddings to recommend diverse articles for users to read.	
	IATA NDC Hackathon 1st Place	05/2016
	Won 5000€ and an incubation program for our app that predicts future travel reimbursements.	
	Cisco DevNet Coding Camp 1st Place	02/2016
	Won 5000€ against 25 teams for an app that integrates with sensors to create and share smart home recipes.	
Ayuda Berlin Code Fight 3rd Place	11/2015	
Competitive programming tournament in a boxing ring. Missed the finals by 3 seconds.		
Google Science Fair Regional Finalist (1 of 30 worldwide)	08/2012	
Invented an activation propagation algorithm to generate word associations from Wikipedia.		
Mathematical Olympiad Germany	2008–2012	
Won at regional and state levels. Invited into Germany's competitive math national team.		