

MICHEL BAKKER

Assistant Professor, Massachusetts Institute of Technology
Senior Research Scientist, Google DeepMind

 bakker@mit.edu
 miba.dev
 [Google Scholar](https://scholar.google.com/citations?user=0000-0002-1234-5678&hl=en)

ACADEMIC APPOINTMENT

9/2024–present **Assistant Professor**
Massachusetts Institute of Technology, Cambridge, MA, USA
Mitsui Career Development Professor in Contemporary Technology, Sloan School of Management
Core faculty, MIT Laboratory for Information & Decision Systems
Affiliated Faculty, MIT Institute for Data, Systems, and Society

EDUCATION

2017–2020 **Ph.D. Computer Science**
Massachusetts Institute of Technology, Cambridge, MA, USA
MIT-IBM Watson AI Lab Fellowship

2017–2019 **M.Sc. Computer Science**
Massachusetts Institute of Technology, Cambridge, MA, USA
Irwin Mark Jacobs and Joan Klein Jacobs Presidential Fellowship

2012–2015 **M.Sc. Physics**
Delft University of Technology, Delft, Netherlands
Cum Laude, Top 5% cumulative GPA, Casimir Honors Program
One-year leave of absence (2013–2014) for entrepreneurial role in Myanmar

2008–2012 **B.Sc. Physics**
Delft University of Technology, Delft, Netherlands

RESEARCH EXPERIENCE

10/2022–present **Senior Research Scientist**
2/2021–10/2022 **Research Scientist**
Google DeepMind, London, United Kingdom
Focus areas: AI Safety, Large language models

1/2017–12/2020 **Graduate Research Assistant**
MIT Media Lab, Cambridge, MA, USA
Advisor: Prof. Alex 'Sandy' Pentland
Thesis: Algorithmic fairness in sequential decision making

8/2019–12/2019 **Research Intern**
Google DeepMind, London, United Kingdom
Supervisors: William S. Isaac, Joel Z. Leibo & Edward Hughes
Project: Modelling cooperation in network games with spatio-temporal complexity

9/2015–12/2015 **Research Intern, Quantum Computing**
IBM Research, Yorktown Heights, NY, USA
Supervisors: Jerry M. Chow & Jay M. Gambetta
Project: Towards faster high-fidelity single-qubit gates in superconducting qubits

12/2014–7/2015 **Master Thesis**
QuTech, Delft University of Technology, Delft, Netherlands
Advisors: Prof. Ronald Hanson & Prof. Tim H. Taminiau
Thesis: One-second coherence for an electron coupled to a multi-qubit nuclear environment

9/2014–11/2014 **Research Intern**
Institute for Quantum Information, RWTH Aachen, Aachen, Germany
Supervisor: Prof. David P. DiVincenzo
Project: Validity of the single-particle description for multi-electron quantum dots

7/2012–12/2012 **Bachelor Thesis**
CERN, Geneva, Switzerland & Nikhef, Amsterdam, Netherlands
Advisor: Prof. Harry van der Graaf
Thesis: The ‘e-brane’, an electron emission membrane for Trixy, a fast particle tracker

NON-RESEARCH WORK EXPERIENCE

2/2025–present **Co-Founder, AI Plan NL**, Amsterdam, Netherlands
Co-founded and provide strategic direction for AI Plan NL (www.aiplan.nl), a non-profit organization working to ensure the Netherlands achieves a leadership position in AI while proactively managing its risks. Together with Dutch AI entrepreneur Jelle Prins, I have raised funding to hire two part-time policy staff. Our work has resulted in three op-eds in leading Dutch newspapers, has shaped the platforms of three major political parties, and informed policy at the Ministries of Foreign Affairs, Economic Affairs, and General Affairs. We are currently co-organizing a Dutch AI summit.

12/2015–12/2016 **Co-Founder, Bloomon**, London, United Kingdom
Co-founded e-commerce flower startup Bloomon in the United Kingdom but left to start my PhD at MIT. Bloomon raised €24 million and was later acquired by UK-based Bloom & Wild. I led the teams in Bloomon’s London office and the business intelligence team in Amsterdam.

7/2013–8/2014 **Co-Founder, Lamudi Myanmar**, Yangon, Myanmar
Built and led a team of 21 full-time employees of a Myanmar online real estate startup. The startup connects real-estate agents with buyers and renters online. I was responsible for the management and expansion of all operations and business activities. The startup was acquired by Alibaba Group.

TEACHING EXPERIENCE

Fall 2024 &
Spring 2025 **Generative AI for Managers, MIT Sloan School of Management**
Co-taught with Prof. John Horton, this course introduces technically inclined MBA students to generative AI and its societal implications. I led all the technical lectures and the lectures on responsible AI and societal implications of AI. Enrollment of 100 students per term from which I received an average teaching evaluation of 5.9/7 for Fall 2024 and 6.4/7 for Spring 2025.

Spring 2025 **Generative AI Action Learning Lab, Massachusetts Institute of Technology**
Co-developed and taught with Prof. John Horton, this MIT-wide technical course connects student teams with external organizations to work on real-world technical generative AI projects with external organizations. The course combines hands-on technical work with guest lectures and a curriculum covering both the technical foundations and responsibility challenges of generative AI. Enrollment of 120 students from which I received an average teaching evaluation of 6.5/7.

2/2018–9/2020 **Sandbox Mentor, MIT Sandbox**, Cambridge, MA, USA
Mentoring students as they explore opportunities to grow their research projects into start-ups. MIT Sandbox provides equity-free seed funding and mentorship for students.

7/2011–8/2013 **Student Assistant, PWS TU Delft**, Delft University of Technology, Netherlands
Provided in-person and online mentoring for high school students to support their final year research project. I also developed and taught workshops on topics like superconductivity and particle physics.

RESEARCH ADVISING

POSTDOCTORAL SUPERVISION

2025-present **Sally Dong** PhD '24 University of Washington Computer Science

GRADUATE ADVISING

2025-present **Xiaoyang Cao** MSc '26, MIT Technology and Policy Program

2025-present **Yanchen Liu** PhD '29, MIT Social and Engineering Systems

2025-present **Sulekha Kishore** PhD '29, MIT Social and Engineering Systems

2025-present **Prerna Ravi** PhD '27, MIT Computer Science (co-advised with Hal Abelson)

2024-present **Haiwen Li** PhD '27, MIT Social and Engineering Systems

2024-present **Soham De** PhD '28, University of Washington Information Science (committee, advised by Martin Saveski)

2024-present **Suyash Fulay** PhD '25 MIT Media Arts and Sciences (co-advised with Deb Roy)

2024-present **Tobin South** PhD '25 MIT Media Arts and Sciences (committee, advised by Alex Pentland)

2024-present **Robin Na** PhD '27, MIT Management Science (committee, advised by Abdullah Almaatouq)

2024-present **Elinor Poole-Dayan** M.Sc.'25 MIT Media Arts and Sciences (co-advised with Deb Roy, now a predoctoral student in my lab)

2024-present **Michael Wong** M.Sc.'25 MIT Media Arts and Sciences (committee, advised by Deb Roy)

UNDERGRADUATE ADVISING & VISITING STUDENTS

Summer 2025 **Jiayi Wu**, B.Sc. '28, Brown University

Summer 2025 **Sercan Demir**, PhD '27, University of Cologne

Summer 2020 **Humberto Riverón Valdés**, B.Sc. '20, Massachusetts Institute of Technology

Summer 2020 **Duy Patrick Tu**, M.Sc. '20, Ludwig Maximilian University of Munich

Fall 2018 **Madelon Hulsebos**, M.Sc. '19, University of Amsterdam (co-mentored with Kevin Hu)

Summer 2018 **Patricia Lu**, B.Sc. '19, Massachusetts Institute of Technology

Summer 2018 **Keis Bejgo**, B.Sc. '19, Massachusetts Institute of Technology

Summer 2018 **Daoud Piracha**, B.Sc. '19, McGill University

Spring 2018 **Stephen Li**, B.Sc. '18, Massachusetts Institute of Technology (co-mentored with Kevin Hu)

PUBLICATIONS

**authors contributed equally*

JOURNAL ARTICLES

Jun 2025 **Haiwen Li**, **Soham De**, **Manon Revel**, **Andreas Haupt**, **Brad Miller**, **Keith Coleman**, **Jay Baxter**, **Martin Saveski**, **Michiel A. Bakker**, *Scaling Human Judgment in Community Notes with LLMs*, *Journal of Online Trust and Safety*, 3(1)

Oct 2025 **Christopher Summerfield**, **Lisa P. Argyle**, **Michiel A. Bakker**, **Teddy Collins**, **Esin Durmus**, **Tyna Eloundou**, **Iason Gabriel**, **Deep Ganguli**, **Kobi Hackenburg**, **Gillian K. Hadfield**, **Luke Hewitt**, **Saffron Huang**, **Hélène Landemore**, **Nahema Marchal**, **Aviv Ovadya**, **Ariel Procaccia**, **Mathias Risse**, **Bruce Schneier**, **Elizabeth Seger**, **Divya Siddarth**, **Henrik Skaug Sætra**, **Michael Henry Tessler**, **Matthew Botvinick**, *The impact of advanced AI systems on democracy*, *Nature Human Behaviour*, 1–11, Nature Publishing Group UK

Nov 2024 Michael Henry Tessler*, **Michiel A. Bakker***, Daniel Jarrett, Hannah Sheahan, Martin J. Chadwick, Raphael Koster, Georgina Evans, Lucy Campbell-Gillingham, Tantum Collins, David C. Parkes, Matthew Botvinick, Christopher Summerfield, *AI can help humans find common ground in democratic deliberation*, Science, 386 (6719), eadq2852

Oct 2024 Jason W. Burton, Ezequiel Lopez-Lopez, Shahar Hechtlinger, Zoe Rahwan, Samuel Aeschbach, **Michiel A. Bakker**, Joshua A. Becker, Aleks Berditchevskiaia, Julian Berger, Levin Brinkmann, Lucie Flek, Stefan M. Herzog, Saffron Huang, Sayash Kapoor, Arvind Narayanan, Anne-Marie Nussberger, Taha Yasseri, Pietro Nickl, Abdulla Almaatouq, Ulrike Hahn, Ralf H. J. M. Kurvers, Susan Leavy, Iyad Rahwan, Divya Siddarth, Alice Siu, Anita W. Woolley, Dirk U. Wulff, Ralph Hertwig, *How large language models can reshape collective intelligence*, Nature Human Behaviour, 1–13

Jun 2024 Edgar A. Duéñez-Guzmán, Ramona Comanescu, Yiran Mao, Kevin R. McKee, Ben Coppin, Suzanne Sadedin, Silvia Chiappa, Alexander S. Vezhnevets, **Michiel A. Bakker**, Yoram Bachrach, William Isaac, Karl Tuyls, Joel Z. Leibo, *Perceptual interventions ameliorate statistical discrimination in learning agents*, Proceedings of the National Academy of Sciences, 122 (25) e2319933121

Sep 2023 Kevin McKee, Andrea Tacchetti, **Michiel A. Bakker**, Jan Balaguer, Lucy Campbell-Gillingham, Richard Everett, Matthew M. Botvinick *Scaffolding cooperation in human groups with deep reinforcement learning*, Nature Human Behavior, 7, 1787–1796

Jun 2022 Alberto Aleta, David Martín-Corral, **Michiel A. Bakker**, Ana Pastore y Piontti, Marco Ajelli, María Litvinova, Matteo Chinazzi, Natalie E. Dean, M. Elizabeth Halloran, Ira M. Longini Jr, Alex Pentland, Alessandro Vespignani, Yamir Moreno, Esteban Moro, *Quantifying the importance and location of SARS-CoV-2 transmission events in large metropolitan areas*, Proceedings of the National Academy of Sciences, 119 (26) e2112182119

Dec 2019 Mohamed H. Abobeih, Joe Randall, Conor E. Bradley, Hans P. Bartling, **Michiel A. Bakker**, Maarten J. Degen, Mark Markham, Daniel J. Twitchen, Tim H. Taminiau, *Atomic-scale imaging of a 27-nuclear-spin cluster using a quantum sensor*, Nature, 576, pages411–415

Sep 2019 Conor E. Bradley, Joe Randall, Mohamed H. Abobeih, Remon C. Berrevoets, Maarten J. Degen, **Michiel A. Bakker**, Mark Markham, Daniel J. Twitchen, Tim H. Taminiau, *A 10-qubit solid-state spin register with quantum memory up to one minute*, Physical Review X, 9, 031045

Jun 2018 Mohamed H. Abobeih, Julia Cramer, **Michiel A. Bakker**, Norbert Kalb, Daniel J. Twitchen, Matthew Markham, Tim H. Taminiau, *One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment*, Nature Communications, 9 2552

Apr 2015 **Michiel A. Bakker***, Sebastian Mehl*, Tuukka Hiltunen, David P. DiVincenzo, *Validity of the single-particle description and noise resilience for multielectron quantum dots*, Physical Review B, 91, 155425

Jan 2013 Harry van der Graaf, **Michiel A. Bakker**, Hong W. Chan, Eduardo Charbon, Fabio Santagata, Pasqualina M. Sarro, Dennis R. Schaar, *Tipsy single soft photon detector and Trixy ultrafast tracking detector*, Journal of Instrumentation, 8 C01036

CONFERENCE PAPERS

Feb 2026 Suyash Fulay, Jocelyn Zhu, **Michiel A. Bakker**, *From Delegates to Trustees: How Optimizing for Long-Term Interests Shapes Bias and Alignment in LLM*, Oral, The International Association for Safe & Ethical AI (IASEAI) 2026

Nov 2025 Taylor Sorensen, Pushkar Mishra, Roma Patel, Michael Henry Tessler, **Michiel A. Bakker**, Georgina Evans, Iason Gabriel, Noah Goodman, Verena Rieser, [Value Profiles for Encoding Human Variation](#), Empirical Methods in Natural Language Processing (EMNLP 2025)

Aug 2025 Michael Henry Tessler, Georgina Evans, **Michiel A. Bakker**, Iason Gabriel, Sophie Bridgers, Rishabh Jain, Raphael Koster, Verena Rieser, Anca Dragan, Matthew Botvinick, Christopher Summerfield, [Can AI Mediation Improve Democratic Deliberation?](#), Artificial Intelligence and Democratic Freedoms, Knight First Amendment Institute

Jul 2025 Aviv Ovadya, Kyle Redman, Luke Thorburn, Quan Ze Chen, Oliver Smith, Flynn Devine, Andrew Konya, Smitha Milli, Manon Revel, Kevin Feng, Amy X Zhang, Bilva Chandra, **Michiel A. Bakker**, Atoosa Kasirzadeh, [Democratic AI is Possible. The Democracy Levels Framework Shows How It Might Work.](#), Forty-second International Conference on Machine Learning (ICML)

Jun 2025 Andrew Konya, Luke Thorburn, Wasim Almasri, Oded Adomi Leshem, Ariel D. Procaccia, Lisa Schirch, **Michiel A. Bakker**, [Using Collective Dialogues and AI to Find Common Ground Between Israeli and Palestinian Peacebuilders](#), ACM Conference on Fairness, Accountability, and Transparency (FAccT)

Apr 2025 Soham De, **Michiel A. Bakker**, Jay Baxter, Martin Saveski, [Supernotes: Driving Consensus in Crowd-Sourced Fact-Checking](#), The 2025 ACM Web Conference (WWW)

Dec 2022 **Michiel A. Bakker***, Martin J. Chadwick*, Hannah R. Sheahan*, Michael Henry Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matthew M. Botvinick, Christopher Summerfield [Fine-tuning language models to find agreement among humans with diverse preferences](#), Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)

Jul 2021 **Michiel A. Bakker**, Duy Patrick Tu, Krishna P. Gummadi, Alex Pentland, Kush R. Varshney, Adrian Weller, [Beyond Reasonable Doubt: Improving Fairness in Budget-Constrained Decision Making using Confidence Thresholds](#), Fifth AAAI/ACM Conference on AI, Ethics, and Society (AIES)

May 2021 **Michiel A. Bakker***, Richard Everett*, Laura Weidinger, Iason Gabriel, William S. Isaac, Joel Z. Leibo, Edward Hughes [Modelling Cooperation in Network Games with Spatio-Temporal Complexity](#), 20th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)

Jan 2020 Alejandro Noriega-Campero, Bernardo Garcia-Bulle-Bueno, Luis Fernando Cantu, **Michiel A. Bakker**, Alex Pentland, [Algorithmic Targeting of Social Policies: Fairness, Accuracy, and Distributed Governance](#), ACM 2020 Conference on Fairness, Accountability and Transparency (FAT*)

Jan 2019 Alejandro Noriega-Campero*, **Michiel A. Bakker***, Bernardo Garcia-Bulle, Alex Pentland [Active Fairness in Algorithmic Decision Making](#), Third ACM Conference on AI, Ethics, and Society (AIES)

Jul 2019 Madelon Hulsebos*, Kevin Z. Hu*, **Michiel A. Bakker**, Emanuel Zgraggen, Tim Kraska, Çagatay Demiralp, César A Hidalgo, [Sherlock: A Deep Learning Approach to Semantic Data Type Detection](#), ACM Conference on Knowledge Discovery and Data Mining (KDD)

May 2019 Kevin Z. Hu, **Michiel A. Bakker**, Stephen Li, Tim Kraska, César A. Hidalgo [VizML: A Machine Learning Approach to Visualization Recommendation](#), ACM Conference on Human Factors in Computing Systems (CHI)

May 2019 Kevin Z. Hu, Neil Gaikwad, Madelon Hulsebos, **Michiel A. Bakker**, Emanuel Zgraggen, César A Hidalgo, Tim Kraska, Guoliang Li, Arvind Satyanarayanan, Cagatay Demiralp, [VizNet: Towards A Large-Scale Visualization Learning and Benchmarking Repository](#), ACM Conference on Human Factors in Computing Systems (CHI)

PRE-PRINTS AND UNDER REVIEW

Dec 2025 Elinor Poole-Dayan, Jiayi Wu, Taylor Sorensen, Jiaxin Pei, **Michiel A. Bakker**, *Benchmarking Overton Pluralism in LLMs*, arXiv preprint arXiv:2512.01351

Oct 2025 Shenzhe Zhu, Shu Yang, **Michiel A. Bakker**, Alex Pentland, Jiaxin Pei, *Can AI Truly Represent Your Voice in Deliberations? A Comprehensive Study of Large-Scale Opinion Aggregation with LLMs*, arXiv preprint arXiv:2510.05154

Jul 2025 Gheorghe Comanici, Eric Bieber, ..., **Michiel A. Bakker**, ..., Yin Li, Chris Hahn (3289 authors ordered randomly), *Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities*, arXiv preprint arXiv:2507.06261

Jul 2025 Anand V. Shah, Tobin South, Talfan Evans, Hannah Rose Kirk, Andrew Trask, E. Glen Weyl, **Michiel A. Bakker**, *Robust AI Personalization Will Require a Human Context Protocol*, Preprint

Jul 2025 Ryan Lowe, Joe Edelman, Tan Zhi-Xuan, Oliver Klingefjord, Ellie Hain, Vincent Wang, Atrisha Sarkar, **Michiel A. Bakker**, Fazl Barez, Matija Franklin, Andreas Haupt, Jobst Heitzig, Wesley H Holliday, Julian Jara-Ettinger, Atoosa Kasirzadeh, Ryan Othniel Kearns, James Ravi Kirkpatrick, Andrew Koh, Joel Lehman, Sydney Levine, Manon Revel, Ivan Vendrov, *Full-Stack Alignment: Co-Aligning AI and Institutions with Thicker Models of Value*, Preprint

Feb 2025 Siddarth Srinivasan, Ezra Karger, **Michiel A. Bakker**, Yiling Chen, *Tell Me Why: Incentivizing Explanations*, arXiv preprint arXiv:2502.13410

Dec 2024 Beth Goldberg, Diana Acosta-Navas, **Michiel A. Bakker**, Ian Beacock, Matt Botvinick, Prateek Buch, Renée DiResta, Nandika Donthi, Nathanael Fast, Ravi Iyer, Zaria Jalan, Andrew Konya, Grace Kwak Danciu, Hélène Landemore, Alice Marwick, Carl Miller, Aviv Ovadya, Emily Saltz, Lisa Schirch, Dalit Shalom, Divya Siddarth, Felix Sieker, Christopher Small, Jonathan Stray, Audrey Tang, Michael Henry Tessler, Amy Zhang, *AI and the Future of Digital Public Squares*, arXiv preprint arXiv:2412.09988

Oct 2024 Christopher Summerfield, Lisa Argyle, **Michiel A. Bakker**, Teddy Collins, Esin Durmus, Tyna Eloundou, Iason Gabriel, Deep Ganguli, Kobi Hackenburg, Gillian Hadfield, Luke Hewitt, Saffron Huang, Helene Landemore, Nahema Marchal, Aviv Ovadya, Ariel Procaccia, Mathias Risse, Bruce Schneier, Elizabeth Seger, Divya Siddarth, Henrik Skaug Sætra, Michael Henry Tessler, Matthew Botvinick, *How will advanced AI systems impact democracy?*, arXiv preprint arXiv:2409.06729

REVIEWING AND AREA CHAIR

2025–2026 Program Committee AI for Civic Discourse, Laude Institute

2020–2025 Conference on Neural Information Processing Systems (NeurIPS)

2023–2024 ACM Conference on Human Factors in Computing Systems (CHI)

2020–2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT)

2021 AAAI /ACM Conference on AI, Ethics, and Society Conference (AIES)

INVITED TALKS

AI-mediated systems for collective decision making and knowledge generation

Dec 2025 International Workshop on Re-imaging Democracy, New Haven, CT

Dec 2025 NeurIPS Workshop on Post-AGI Futures

	AI can help humans find common ground in democratic deliberation
Oct 2025	INFORMS, Atlanta, GA
Jun 2025	9th Annual Center for Human-Compatible AI Workshop, Asilomar, CA
Jun 2025	AI, Science and Innovation Seminar, Northwestern University
Apr 2025	EconCS Seminar, Harvard University, Cambridge, MA, USA
Mar 2025	Program on Negotiation AI Summit, Cambridge, MA, USA
Feb 2025	Conference of the International Association for Safe and Ethical Artificial Intelligence (IASEAI '25), Paris, France
Feb 2025	Rebooting Democracy in the Age of AI lecture, Northeastern University, MA, USA
Jan 2025	Newspeak House, London, UK
Dec 2024	International Workshop on Reimagining Democracy 2024, Washington, DC, USA
May 2024	Conference on AI's impact on Society, Media & Democracy, University of Amsterdam
	AI, deliberation and alignment
June 2024	Keynote (with Anca Dragan), 8th Annual Center for Human-Compatible AI Workshop, Asilomar, CA
	Fine-tuning language models to find agreement among humans with diverse preferences
Apr 2023	Center for Humans and Machines Seminar, Max Planck Institute for Human Development
Mar 2023	Amsterdam Machine Learning Lab, University of Amsterdam
Mar 2023	Institute for Logic, Language and Computation, University of Amsterdam
Feb 2023	Sloan School of Management, Massachusetts Institute of Technology
Feb 2023	Generative AI for Constructive Communication Course, Massachusetts Institute of Technology
Dec 2022	OpenAI Alignment Team Meeting
Dec 2022	MIT Center for Collective Intelligence Seminar, Massachusetts Institute of Technology
Dec 2022	Stanford NLP Seminar, Standford University
	Fairness in Budget-Constrained Decision Making
Nov 2019	Ethics DeepDive Team Meeting, DeepMind
Apr 2019	Imagination in Action, MIT Media Lab
Nov 2018	Data-Pop Alliance
	A Large-Scale Study of Social Integration of Syrian Refugees in Turkey
Jan 2019	Imagination in Action, World Economic Forum
Jan 2019	Boğaziçi University
	Quantum Computing with Diamond
Dec 2015	IBM Q Conference New York
Sep 2015	QuTech Seminar, Delft University of Technology
	Towards Noise-free Single-triplet Qubits
Jul 2015	QuTech Seminar, Delft University of Technology