

MICHIEL BAKKER

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

✉ bakker@mit.edu
🌐 miba.dev

ACADEMIC APPOINTMENT

- 9/2024–present **Assistant Professor**
Massachusetts Institute of Technology, Cambridge, MA, USA
- 9/2024–present Assistant Professor of Information Technology, MIT Sloan School of Management
- 9/2024–present Affiliated Faculty, MIT Institute for Data, Systems, and Society
- 10/2024–present Affiliated Faculty, MIT Center for Constructive Communication, MIT Media Lab

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. Applied Physics**
Delft University of Technology, Delft, Netherlands
Cum Laude, Top 5% cumulative GPA, Casimir Honors Program
- 2013–2014 One-year break for entrepreneurial experience in Myanmar
- 2008–2012 **B.Sc. Applied 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 **PhD Thesis**
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

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 **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. Received an average teaching evaluation of 9.1/10.

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 TUDelft**, 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

2024-present **Haiwen Li** PhD '27, Massachusetts Institute of Technology (co-advised with Sinan Aral)
2024-present **Suyash Fulay** PhD '25 Massachusetts Institute of Technology (co-advised with Deb Roy)
2024-present **Tobin South** PhD '25 Massachusetts Institute of Technology (co-advised with Alex Pentland)
2024-present **Elinor Poole-Dayam** M.Sc.'25 Massachusetts Institute of Technology (co-advised with Deb Roy)
2024-present **Michael Wong** M.Sc.'25 Massachusetts Institute of Technology (co-advised with Deb Roy)
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-advised with Kevin Zeng 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-advised with Kevin Zeng Hu)

PUBLICATIONS

**authors contributed equally*

JOURNAL ARTICLES

- 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 Berdichevskaia, Julian Berger, Levin Brinkmann, Lucie Flek, Stefan M. Herzog, Saffron Huang, Sayash Kapoor, Arvind Narayanan, Anne-Marie Nussberger, Taha Yasseri, Pietro Nickl, Abdullah 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
- July 2024 Edgar A. Duéñez-Guzmán, Kevin R. McKee, Yiran Mao, Ben Coppin, Silvia Chiappa, Alexander Sasha Vezhnevets, **Michiel A. Bakker**, Yoram Bachrach, Suzanne Sadedin, William Isaac, Karl Tuyls, Joel Z. Leibo, *Statistical discrimination in learning agents*, Proceedings of the National Academy of Sciences (PNAS), in press
- 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, Maria 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. Schaart, *Tipsy single soft photon detector and Trixy ultrafast tracking detector*, Journal of Instrumentation, 8 C01036

CONFERENCE PAPERS

- 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) [26% acceptance rate]
- 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) [25% acceptance rate]
- 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) [29% acceptance rate]
- 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*) [24% acceptance rate]
- Jan 2019 Alejandro Noriega-Campero*, **Michiel A. Bakker***, Bernardo Garcia-Bulle, Alex Pentland *Active Fairness in Algorithmic Decision Making*, Third AAAI/ACM Conference on AI, Ethics, and Society (AIES) [Selected for oral presentation, 15% acceptance rate]
- Jul 2019 Madelon Hulsebos*, Kevin Z. Hu*, **Michiel A. Bakker**, Emanuel Zraggen, Tim Kraska, Çağatay Demiralp, César A Hidalgo, *Sherlock: A Deep Learning Approach to Semantic Data Type Detection*, ACM Conference on Knowledge Discovery and Data Mining (KDD) [14% acceptance rate]
- 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) [23.8 % acceptance rate]
- May 2019 Kevin Z. Hu, Neil Gaikwad, Madelon Hulsebos, **Michiel A. Bakker**, Emanuel Zraggen, César A Hidalgo, Tim Kraska, Guoliang Li, Arvind Satyanarayan, Çağatay Demiralp, *VizNet: Towards A Large-Scale Visualization Learning and Benchmarking Repository*, ACM Conference on Human Factors in Computing Systems (CHI) [23.8 % acceptance rate]

WORKSHOP PAPERS (refereed)

- Sep 2023 Daniel Jarrett, Miruna Pislari, **Michiel A. Bakker**, Michael Henry Tessler, Raphael Koster, Jan Balaguer, Romuald Elie, Christopher Summerfield, Andrea Tacchetti, *Language agents as digital representatives in collective decision-making*, NeurIPS 2023 Foundation Models for Decision Making Workshop
- Feb 2020 **Michiel A. Bakker**, Humberto Riverón Valdés, Duy Patrick Tu, Krishna P. Gummadi, Kush R. Varshney, Adrian Weller, Alex Pentland, *Fair Enough: Improving Fairness in Budget-Constrained Decision Making Using Confidence Thresholds* SafeAI, Artificial Intelligence Safety workshop at AAAI 2019 [Selected for oral presentation]
- Oct 2019 **Michiel A. Bakker**, Duy Patrick Tu, Humberto Riverón Valdés, Krishna P. Gummadi, Kush R. Varshney, Adrian Weller, Alex Pentland, *DADI: Dynamic Discovery of Fair Information with Adversarial Reinforcement Learning* Human-Centric Machine Learning workshop at NeurIPS 2019

- Dec 2019 Alejandro Noriega-Campero, Bernardo Garcia Bulle-Bueno, Luis Fernando Cantu, **Michiel A. Bakker**, Luis Tejerina, Alex Pentland, *Algorithmic Targeting of Social Policies: Accuracy & Fairness*, Machine Learning for the Developing World workshop at NeurIPS 2019
- Dec 2019 **Michiel A. Bakker**, Humberto Riverón Valdés, Duy Patrick Tu, Krishna P. Gummadi, Kush R. Varshney, Adrian Weller, Alex Pentland, *On Fairness in Budget-Constrained Decision Making* Explainable AI for Accountability, Fairness, and Transparency workshop at KDD 2019 [Best paper award, selected for oral presentation]
- Jul 2018 Bjarke Felbo*, **Michiel A. Bakker***, Abhimanyu Dubey, Sadhika S. Malladi, Alex Pentland, Iyad Rahwan, *Prediction Propagation for Domain Adaptation*, Towards learning with limited labels workshop at ICML 2018

PRE-PRINTS AND UNDER REVIEW

- 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, H el ene 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
- Nov 2024 Soham De, **Michiel A. Bakker**, Jay Baxter, Martin Saveski, *Supernotes: Driving Consensus in Crowd-Sourced Fact-Checking*, arXiv preprint arXiv:2411.06116
- 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 S etra, Michael Henry Tessler, Matthew Botvinick, *How will advanced AI systems impact democracy?*, arXiv preprint arXiv:2409.06729

REVIEWING

- 2023–2024 ACM Conference on Human Factors in Computing Systems (CHI)
- 2020–2024 Conference on Neural Information Processing Systems (NeurIPS)
- 2020–2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT)
- 2021 AAAI /ACM Conference on AI, Ethics, and Society Conference (AIES)

INVITED TALKS

- AI can help humans find common ground in democratic deliberation**
- Feb 2025 Rebooting Democracy in the Age of AI lecture, Northeastern University
- Jan 2025 Newspeak House, London, UK
- Dec 2024 International Workshop on Reimagining Democracy 2024, Washington DC
- 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, Stanford 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