

CURRICULUM VITA

Melanie Mitchell

Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501
Email: mm@santafe.edu
Web site: melaniemitchell.me

RESEARCH EXPERIENCE AND INTERESTS

Artificial intelligence, machine learning, biologically inspired computing, cognitive science, and complex systems.

GRADUATE EDUCATION

Ph.D. in Computer Science, 1990, University of Michigan.

Dissertation advisor: Douglas R. Hofstadter

Dissertation title: *Copycat: A Computer Model of High-Level Perception and Conceptual Slippage in Analogy-Making.*

PROFESSIONAL EMPLOYMENT

Professor, Santa Fe Institute, 2020–present.

Professor, Portland State University, Department of Computer Science, 2004–2020.

Professor, July-September 2004; Associate Professor, August 2002-June 2004, OGI School of Science and Engineering, Oregon Health & Science University, Department of Computer Science and Engineering.

Staff Member, Santa Fe Institute, 2000–2002.

Technical Staff Member, Los Alamos National Laboratory, Biophysics Group, 1999–2000.

Research Professor and Director of the Adaptive Computation Program, Santa Fe Institute, 1992–1999.

Postdoctoral Scholar, University of Michigan, 1990–1992.

HONORS AND AWARDS

Eric & Wendy Schmidt Award for Excellence in Science Communications, National Academy of Sciences, 2025.

Senior Scientific Award, Complex Systems Society, 2023.

Artificial Intelligence: A Guide for Thinking Humans named as one of the **five best books on AI** by the *New York Times* and the *Wall Street Journal*.

Artificial Intelligence: A Guide for Thinking Humans shortlisted (one of five) for the **2023 Cosmos Prize for Scientific Writing**.

Distinguished Cognitive Scientist Award, Cognitive and Information Sciences Faculty, UC Merced, 2022.

Herbert A. Simon Award, International Conference on Complex Systems, 2020.

Phi Beta Kappa Society 2010 Science Book Award for *Complexity: A Guided Tour*.

Complexity: A Guided Tour named by Amazon.com as #3 of the **10 best science books of 2009**.

Complexity: A Guided Tour longlisted (one of 12) for the **2010 Royal Society Science Book Prize**.

Ulam Memorial Lectureship, Santa Fe Institute, 1997 (Annual lectureship for three public lectures on Complex Systems).

Fellowship, Michigan Society of Fellows, University of Michigan, 1990–1993 (a three-year postdoctoral fellowship at the University of Michigan).

Regents' Fellowship, University of Michigan, 1984–1988 (a four-year fellowship for graduate studies at the University of Michigan).

GRANTS

The AI Research Institute on Interaction for AI Assistants (ARIA). Ellie Pavlick, PI. NSF AI Institute Program, 2025–2030. (Approximately \$800,000 to Melanie Mitchell.)

Principled Approaches to Research and Development of Innovative Generative Models. Sivasankaran Rajamanickam, PI. Sandia National Laboratory, 2024–2026. (\$603,881 to Melanie Mitchell).

Building Diverse Intelligences through Compositionality and Mechanism Design. J. Foster, PI. Templeton World Charity Foundation, 2022–2026 (\$280,929 to Melanie Mitchell).

EAGER: Developing data and evaluation methods to assess the generality and robustness of AI systems for abstraction and analogy-making. M. Mitchell, PI. National Science Foundation, 2021–2023 (\$199,661).

AI Institute: Planning: Foundations of Intelligence in Natural and Artificial Systems. M. Mitchell, PI; M. Moses, Co-PI. National Science Foundation, 2020–2022 (\$499,918).

Workshop: Artificial Intelligence and the “Barrier of Meaning”. M. Mitchell, PI. National Science Foundation, 2018. (National Science Foundation \$20,000; Artificial Intelligence Journal, \$12,000).

Visual Situation Recognition: An Integration of Deep Networks and Analogy-Making. M. Mitchell, PI. National Science Foundation, 2014–2018. (\$450,000; \$26,000 REU Supplements).

The Principles of Complexity. Grant to the Santa Fe Institute from the John Templeton Foundation, 2011–2013. M. Mitchell was co-PI of *Complexity Explorer* part of this grant. (\$5M total to SFI).

A Scalable Architecture for Image Interpretation, M. Mitchell and G. Kenyon, Co-PIs. National Science Foundation, 2010–2014. (\$500,000).

Workshop: Shared Organizing Principles in the Biological and Computing Sciences, D. McPheeters, PI; M. Mitchell, J. Wise, and R. Greenspan, Co-PIs. National Science Foundation, 2010. (\$98,570).

Evolving and Understanding Cellular Arrays, M. Mitchell, PI. MARCO Center for Functional Engineered Nano Architectonics (funded by Microelectronics Advanced Research Corporation), 2006–2009 (\$325,000).

Adaptive Image Feature Detection and Image Retrieval with Genetic Algorithms, M. Mitchell, PI. Intel Corporation, 2003–2005 (\$102,996).

Perception and Analogy-Making in Complex Adaptive Systems, M. Mitchell, PI. J. S. McDonnell Foundation, 2002–2007 (\$377,533).

Complex Systems Summer School, R. Goldstein and M. Mitchell, PIs. Office of Naval Research, 2000–2002 (\$294,855).

Computing with Networks of Spiking Neurons, M. Mitchell, PI. Los Alamos National Laboratory (Laboratory Directed Research and Development Grant), 2000 (\$150,000).

1999 Complex Systems Summer School, M. Mitchell and L. Nadel, PIs. Department of Energy, 1999–2000 (\$50,000).

Evolving Cellular Automata with Genetic Algorithms, M. Mitchell, PI; J. P. Crutchfield, Co-PI. National Science Foundation, 1997–1999 (\$224,947).

Rapid Feature Identification using Reconfigurable Computing Technology and Genetic Algorithms, M. Mitchell, Subcontractor. Los Alamos National Laboratory, 1998–1999 (\$41,422).

Automatic Programming of Decentralized Parallel Architectures, M. Mitchell, PI. National Science Foundation, CISE Postdoctoral Associate Program, 1997–1999 (\$46,200, with an equal amount in matching funds from the Santa Fe Institute, to support a postdoctoral fellow).

Major Research Instrumentation Program: A New Computational Infrastructure at the Santa Fe Institute, E. Jen, PI; J. P. Crutchfield, S. Durlauf, C. Langton, and M. Mitchell, Co-PIs. National Science Foundation, 1997–2000 (\$202,377).

Evolving Cellular Automata with Genetic Algorithms, J. P. Crutchfield, PI; M. Mitchell, Co-PI. National Science Foundation, 1994–1996 (\$297,828).

Evolving Cellular Automata with Genetic Algorithms, J. Crutchfield and M. Mitchell, PIs. Department of Energy, 1994–1996 (\$120,000).

Foundations of Genetic Algorithms, S. Forrest and M. Mitchell, PIs. Alfred P. Sloan Foundation, 1993–1994 (\$30,000).

PUBLICATIONS

Books Authored

Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. New York: Farrar, Straus, and Giroux.

Mitchell, M. (2009). *Complexity: A Guided Tour*. New York: Oxford University Press.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press.

Mitchell, M. (1993). *Analogy-Making as Perception: A Computer Model*. Cambridge, MA: MIT Press.

Books Edited

Booker, L., Forrest, S., Mitchell, M., and Riolo, R. L. (2005). *Perspectives on Adaptation in Natural and Artificial Systems*. New York: Oxford University Press.

Belew, R. K. and Mitchell, M. (editors). (1996). *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Reading, MA: Addison-Wesley.

Book Chapters

Mitchell, M. (2024). Large language models. In M. C. Frank & A. Majid (Eds.), *Open Encyclopedia of Cognitive Science*. MIT Press.

Mitchell, M. (2024). Symbols versus cybernetics: Marvin Minsky on the prospects for artificial intelligence. Commentary on M. Minsky, “Steps toward artificial intelligence” (1961). In D. C. Krakauer (editor), *Foundational Papers in Complexity Science, Volume I*. Santa Fe, NM: SFI Press.

Mitchell, M. (2021). Why AI is harder than we think. To appear in C. Craver and C. Klein (editors), *Mind Design 3*.

Ghosh, P., Mitchell, M., Tanyi, J., and Hung, A. (2009). A genetic algorithm-based level-set curve evolution for prostate segmentation on pelvic CT and MRI images. E. Romero and F. Gonzalez (editors), *Biomedical Image Analysis and Machine Learning Technologies: Applications and Techniques*. ICI Global.

Mitchell, M. (2008). Five questions. In C. Gershenson (editor), *Complexity: 5 Questions*. Automatic Press.

- Cenek, M. and Mitchell, M. (2007). Evolving cellular automata. In R. A. Meyers (editor), *Encyclopedia of Complexity and Systems Science*. Berlin: Springer-Verlag.
- Mitchell, M. (2006). Coevolutionary learning with spatially distributed populations. In G. Y. Yen and D. B. Fogel (editors), *Computational Intelligence: Principles and Practice*. New York: IEEE Computational Intelligence Society.
- Crutchfield, J. P., Mitchell, M., and Das, R. (2003). Evolutionary design of collective computation in cellular automata. In J. P. Crutchfield and P. K. Schuster (editors), *Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function*, pp. 361–411. New York: Oxford University Press.
- Mitchell, M. and Newman, M. (2002). Complex systems theory and evolution. In M. Pagel (editor), *Encyclopedia of Evolution*. New York: Oxford University Press.
- Mitchell, M. (2001). Analogy-making as a complex adaptive system. In L. Segel and I. Cohen (editors), *Design Principles for the Immune System and Other Distributed Autonomous Systems*. New York: Oxford University Press.
- Mitchell, M. (1999). Evolutionary computation. In R. Wilson and F. Keil (editors), *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.
- Mitchell, M. (1998). Computation in cellular automata: A selected review. In T. Gramss, S. Bornholdt, M. Gross, M. Mitchell, and T. Pellizzari, *Nonstandard Computation*, pp. 95–140. Weinheim: VCH Verlagsgesellschaft.
- Mitchell, M., Crutchfield, J. P., and Das, R. (1998). Evolving cellular automata to perform computations. In T. Bäck, D. Fogel, and Z. Michalewicz (editors), *Handbook of Evolutionary Computation*. Oxford: Oxford University Press.
- Mitchell, M. and Forrest, S. (1998). Royal Road functions. In T. Bäck, D. Fogel, and Z. Michalewicz (editors), *Handbook of Evolutionary Computation*. Oxford: Oxford University Press.
- Belew, R. K., Mitchell, M., and Ackley, D. H. (1996). Computation and the natural sciences. In R. K. Belew and M. Mitchell (editors), *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Reading, MA: Addison-Wesley.
- Mitchell, M., Crutchfield, J. P., and Hraber, P. T. (1994). Dynamics, computation, and the “edge of chaos”: A re-examination. In G. Cowan, D. Pines, and D. Melzner (editors), *Complexity: Metaphors, Models, and Reality*. Reading, MA: Addison-Wesley.
- Mitchell, M. (1993). Genetic algorithms. In L. Nadel and D. L. Stein (editors), *1992 Lectures in Complex Systems*. Reading, MA: Addison-Wesley.
- Hofstadter, D. R. and Mitchell, M. (1994). The Copycat project: A model of mental fluidity and analogy-making. In K. Holyoak and J. Barnden (editors), *Advances in Connectionist and Neural Computation Theory, Volume 2: Analogical Connections*. Norwood, NJ: Ablex Publishing Corporation. (Adapted as two chapters in D. R. Hofstadter and the Fluid Analogies Research Group, *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York: Basic Books, 1995.)

Journal Publications

Mitchell, M., Foster, J. G., Beger, C., Yi, R., Fu, S., Denton, K. K., and Palmarini, A. B. (2026). Illusions of abstraction in contemporary AI. To appear in *Proceedings of the National Academy of Sciences*.

Mitchell, M. (2026). On evaluating abstraction and analogy in humans and machines. *Current Directions in Psychological Science*.

Johnson, S. G. B., Karimi, A. H., Bengio, Y., Chater, N., Gerstenberg, T., Larson, K., Levine, S., Mitchell, M., Rahwan, I., Schölkopf, B., and Grossmann, I. (2026). Imagining and building wise machines: The centrality of AI metacognition. *Trends in Cognitive Science*.

Stevenson, C. E., Pafford, A., van der Maas, H. L. J., and Mitchell, M. (2025). Can large language models generalize analogy solving like children can? To appear in *Transactions of the Association for Computational Linguistics*.

Krakauer, D. C., Krakauer, J. W., and Mitchell, M. (2025). Large language models and emergence: A complex systems perspective. To appear in *Philosophical Transactions of the Royal Society A*.

Lewis, M. and Mitchell, M. (2025). Evaluating the robustness of analogical reasoning in GPT models. *Transactions on Machine Learning Research*.

Moskvichev, A., Odouard, V. V., and Mitchell, M. (2023). The ConceptARC benchmark: Evaluating understanding and generalization in the ARC domain. *Transactions on Machine Learning Research*.

Burnell, R., Schellaert, W., Burden, J., Ullman, T. D., Martinez-Plumed, F., Tenenbaum, J. B., Rutar, D., Cheke, L. G., Sohl-Dickstein, J., Mitchell, M., Kiela, D., Shanahan, M., Voorhees, E. M., Cohn, A. G., Leibo, J. Z., and Hernandez-Orallo, J. (2023). Rethink reporting of evaluation results in AI. *Science*, 380 (6641), 136-138.

Mitchell, M. and Krakauer, D. C. (2023). The debate over understanding in AI's large language models. *Proceedings of the National Academy of Sciences*, 120 (13).

Shiffren, R. and Mitchell, M. (2023). Probing the psychology of AI models. *Proceedings of the National Academy of Sciences*, 120 (10).

Mitchell, M. (2021). Abstraction and analogy-making in artificial intelligence. *Annals of the New York Academy of Science*, 1505 (1), 79–101.

Mitchell, M. (2020). On crashing the barrier of meaning in AI. *AI Magazine*, 41(2), 86–92.

Mitchell, M. (2019). Artificial intelligence hits the barrier of meaning. *Information*, 10 (2), 51.

Lizier, J. T., Harré, M. S., Mitchell, M., DeDeo, S., Finn, C., Lindgren, K., Lizier, A. L., and Sayama, H. (2018). An interview-based study of pioneering experiences in teaching and learning complex systems in higher education. *Complexity* 1, 1–16.

- Forrest, S. and Mitchell, M. (2016). Adaptive computation: The multidisciplinary legacy of John H. Holland. *Communications of the ACM*, 59 (8), 58–63.
- Ghosh, P., Mitchell, M., Tanyi, J. A., and Hung, A. Y. (2016). Incorporating priors for medical image segmentation using a genetic algorithm. *Neurocomputing*. 195, 181–194.
- Mitchell, M. (2012). Biological computation. *The Computer Journal*, 55, 852–855.
- Ghosh, P., Mitchell, M., and Gold, J. (2010). LSGA: Combining level-sets and genetic algorithms for segmentation. *Evolutionary Intelligence*, 3, 1–11.
- Mitchell, M. (2006). Complex systems: Network thinking. *Artificial Intelligence*, 170 (18), 1194–1212.
- Pagie, L. and Mitchell, M. (2002). A comparison of evolutionary and coevolutionary search. *International Journal of Computational Intelligence and Applications*, 2(1), 53–69.
- Mitchell, M. (2001). Life and evolution in computers. *History and Philosophy of the Life Sciences*, 23, 361–383.
- Jimenez-Morales, F., Crutchfield, J. P., and Mitchell, M. (2001). Evolving two-dimensional cellular automata to perform density classification: A report on work in progress. *Parallel Computing*, 27 (5), 571–585.
- Werfel, J., Mitchell, M., and Crutchfield, J. P. (2000). Resource sharing and coevolution in evolving cellular automata. *IEEE Transactions on Evolutionary Computation*, 4(4), 388–393.
- Mitchell, M. and Taylor, C. E. (1999) Evolutionary computation: An overview. *Annual Review of Ecology and Systematics*, 30, 593–616
- Mitchell, M. (1999). Can evolution explain how the mind works? A review of the evolutionary psychology debates. *Complexity*, 3 (3), 17–24.
- Mitchell, M. (1998). Theories of structure versus theories of change. (Commentary on “The dynamical hypothesis in cognitive science”, by T. van Gelder.) *Behavioral and Brain Sciences*, 21, 645–646.
- van Nimwegen, E., Crutchfield, J. P., and Mitchell, M. (1999). Statistical dynamics of the Royal Road genetic algorithm. *Theoretical Computer Science*, 229 (1), 41–102.
- van Nimwegen, E., Crutchfield, J. P., and Mitchell, M. (1997). Finite populations induce metastability in evolutionary search. *Physics Letters A*, 229 (2), 144–150.
- Crutchfield, J. P., and Mitchell, M. (1995). The evolution of emergent computation. *Proceedings of the National Academy of Sciences*, 92 (23): 10742.
- Mitchell, M. (1995). Genetic algorithms: An overview. *Complexity*, 1 (1) 31–39.
- Mitchell, M., Crutchfield, J. P., and Hraber, P. (1994). Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D*, 75, 361–391.
- Mitchell M., and Forrest, S. (1994). Genetic algorithms and artificial life. *Artificial Life*, 1 (3), 267–289.

Mitchell, M., Hraber, P. T., and Crutchfield, J. P. (1993). Revisiting the edge of chaos: Evolving cellular automata to perform computations. *Complex Systems*, 7, 89–130.

Forrest, S. and Mitchell, M. (1993). What makes a problem hard for a genetic algorithm? Some anomalous results and their explanation. *Machine Learning*, 13, 285-319.

Mitchell, M. and Hofstadter, D. R. (1990). The emergence of understanding in a computer model of concepts and analogy-making. *Physica D*, 42, 322–334.

Seward, F. D. and Mitchell, M. (1981). An X-ray survey of the Small Magellanic Cloud. *Astrophysical Journal*, 243, 736.

Mitchell, M. (1979). Period changes in two W Virginis variables. *Journal of the American Association of Variable Star Observers*, 8 (2).

Conference Publications

Beger, C., Yi, R., Fu, S., Moskvichev, A., and Mitchell, M. (2025). Investigating abstraction capabilities of the o3 model using textual and visual modalities. In *Proceedings of the NeurIPS Multimodal Algorithmic Reasoning Workshop*.

Palmarini, A. B. and Mitchell, M. (2024). Abstract understanding of core-knowledge concepts: Humans vs. LLMs. In *Proceedings of the ICML Workshop on Large Language Models and Cognition*.

Lewis, M. and Mitchell, M. (2024). Using counterfactual tasks to evaluate the generality of analogical reasoning in large language models. In *Proceedings of the Cognitive Science Society Conference*.

Mitchell, M., Palmarini, A. B., and Moskvichev, A. (2023). Comparing humans, GPT-4, and GPT-4V on abstraction and reasoning tasks. In *Proceedings of the LLM-CP Workshop (AAAI-2024)*.

Odouard, V. V. and Mitchell, M. (2022). Evaluating understanding on conceptual abstraction benchmarks. In *Proceedings of the Workshop on AI Evaluation Beyond Metrics (IJCAI 2022)*.

Shanahan, M. and Mitchell, M. (2022). Abstraction for deep reinforcement learning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2022)*.

Springer, J. M., Mitchell, M., and Kenyon, G. T. (2021). A little robustness goes a long way: Leveraging robust features for targeted transfer attacks. In *Proceedings of 35th Conference on Neural Information Processing Systems (NeurIPS 2021)*.

Springer, J. M., Mitchell, M., and Kenyon, G. T. (2021). Uncovering universal features: How adversarial training improves adversarial transferability. Workshop on Adversarial Machine Learning, International Conference on Machine Learning (ICML).

Conser, E., Hahn, K., Watson, C., and Mitchell, M. (2019). Revisiting visual grounding. In *Proceedings of the Workshop on Shortcomings in Vision and Language (SiVL 2019)*. Association for Computational Linguistics.

- Rhodes, A. D., Witte, J., Mitchell, M., and Jedynak, B. (2018). Gaussian processes with context-supported priors for active object localization. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN 2018)*.
- Quinn, M. H., Conser, E., Witte, J. M., and Mitchell, M. (2018). Semantic image retrieval via active grounding of visual situations. In *Proceedings of the 12th International Conference on Semantic Computing*. IEEE
- Rhodes, A. D., Witte, J., Mitchell, M., and Jedynak, B. (2017). Bayesian optimization for refining object proposals. In *Proceedings of the Seventh International Conference on Image Processing Theory, Tools, and Applications (IPTA 2017)*. IEEE.
- Lundquist, S. Y., Mitchell, M., and Kenyon, G. T. (2017). Sparse coding on stereo video for object detection. In *Proceedings of the NIPS 2017 Workshop on Learning with Limited Labeled Data: Weak Supervision and Beyond*.
- Rhodes, A. D., Quinn, M. H., and Mitchell, M. (2017). Fast on-line kernel density estimation for active object localization. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN 2017)*.
- Mitchell, M. (2017). Active interpretation of visual situations. In *AAAI Spring Symposium Series*. AAAI Press.
- Thomure, M. D., Mitchell, M., and Kenyon, G. T. (2013). On the role of shape prototypes in hierarchical models of vision. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN 2017)*.
- Landecker, W., Thomure, M. D., Bettencourt, L. M. A., Mitchell, M., Kenyon, G. T., and Brumby, S. P. (2013). Interpreting individual classifications of hierarchical networks. In *Proceedings of the 2013 Conference on Computational Intelligence and Data Mining (CIDM 2013)*.
- Landecker, W., Thomure, M. D., and Mitchell, M. (2011). Background cues in images classified by hierarchical models. (Abstract.) In *Proceedings of Grand Challenges in Neural Computation II: Neuromimetic Processing and Synthetic Cognition*. Santa Fe, NM.
- Thomure, M. D., Landecker, W., and Mitchell, M. (2011). Random prototypes in hierarchical models of vision. (Abstract.) In *Proceedings of Grand Challenges in Neural Computation II: Neuromimetic Processing and Synthetic Cognition*. Santa Fe, NM.
- Ghosh, P., Mitchell, M., and Gold, J. (2010). Segmentation of thermographic images of hands using a genetic algorithm. In *Proceedings of SPIE*, Vol. 7538, 75380D (2010).
- Marques-Pita, M., Mitchell, M., and Rocha, L. (2008). The role of conceptual structure in designing cellular automata to perform collective computation. In *Proceedings of the Conference on Unconventional Computation*, UC 2008, Springer (Lecture Notes in Computer Science).
- Ghosh, P. and Mitchell, M. (2008). Prostate segmentation on pelvic CT images using a genetic algorithm. In *Proceedings of the International Society for Optical Engineering (SPIE), Conference on Medical Imaging*, February, 2008. SPIE Press.
- Juengling, R. and Mitchell, M. (2007). Combinatorial shape decomposition. In *Proceedings of the Third International Symposium on Visual Computing (ISVC07)*. Springer (Lecture Notes in Computer Science).
- Ghosh, P. and Mitchell, M. (2006). Medical image segmentation with genetic algorithms. In *Proceedings of the*

Genetic and Evolutionary Computation Conference, GECCO-2006, pp. 1171–1178.

Mitchell, M., Thomure, M. D., and Williams, N. L. (2006). The role of space in the success of coevolutionary learning. In L. M. Rocha et al. (editors), *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pp. 118–124. Cambridge, MA: MIT Press.

Williams, N. and Mitchell, M. (2005). Investigating the success of spatial coevolutionary learning. In H. G. Beyer et al. (editors), *Proceedings of the 2005 Genetic and Evolutionary Computation Conference, GECCO-2005*. New York: ACM Press, 523–530.

Mitchell, M. (2005). Self-awareness and control in decentralized systems. In *Working Papers of the AAAI 2005 Spring Symposium on Metacognition in Computation*. Menlo Park, CA: AAAI Press.

Jimenez-Morales, F., Mitchell, M., and Crutchfield, J. P. (2002). Evolving one-dimensional cellular automata to perform a non-trivial collective behavior task: One case study. In P. M. A. Slood, C. J. K. Tan, J. J. Dongarra and A. G. Hoekstra (editors), *Computational Science-ICCS 2002, Part I, Proceedings 2329*, 793-802. Berlin: Springer-Verlag.

Brumby, S. P., Perkins, S. J., Theiler, J., Szymanski, J. J., Bloch, J. J., and Mitchell, M. (1999). Investigation of image feature extraction by a genetic algorithm. In *Proceedings of the International Society for Optical Engineering, Proceedings of SPIE 3812*, 24-31. Bellingham, WA: SPIE Press.

Hordijk, W., Crutchfield, J. P., and Mitchell, M. (1998). Mechanisms of emergent computation in cellular automata. In A. E. Eiben (ed.), *Proceedings of the Fifth International Conference on Parallel Problem Solving From Nature—PPSN V*. New York: Springer.

Jimenez-Morales, F., Crutchfield, J. P., and Mitchell, M. (1998). Evolving two-dimensional cellular automata to perform density classification: A report on work in progress. In S. Bandini, R. Serra, and F. Suggi Liverani (eds.), *Cellular Automata: Research Towards Industry (Proceedings of the Third International Conference on Cellular Automata for Research and Industry)*, 3-14. Springer-Verlag.

Mitchell, M. (1998). A complex-systems perspective on the “computation vs. dynamics” debate in cognitive science. In M. A. Gernsbacher and S. J. Derry (eds.), *Proceedings of the 20th Annual Conference of the Cognitive Science Society—Cogsci98*, 710-715.

Mitchell, M., Crutchfield, J. P., and Das, R. (1996). Evolving cellular automata to perform computations: A review of recent work. In *Proceedings of the First International Conference on Evolutionary Computation and its Applications (EvCA '96)*, 42–55. Moscow, Russia: Russian Academy of Sciences.

Hordijk, W., Crutchfield, J. P., and Mitchell, M. (1996). Embedded particle computation in evolved cellular automata. In *Proceedings of the Conference on Physics and Computation—PhysComp96*, Boston, MA.

Das, R., Crutchfield, J. P., Mitchell, M., and Hanson, J. E. (1995). Evolving globally synchronized cellular automata. In L. J. Eshelman (editor), *Proceedings of the Sixth International Conference on Genetic Algorithms*. San Mateo, CA: Morgan Kaufmann.

Das, R., Mitchell, M., and Crutchfield, J. P. (1994). A genetic algorithm discovers particle-based computation in cellular automata. In Y. Davidor, H.-P. Schwefel, and R. Männer (editors), *Parallel Problem Solving from Nature—PPSN III*. Berlin: Springer-Verlag.

Mitchell, M., Holland, J. H., and Forrest, S. (1994). When will a genetic algorithm outperform hill climbing? In J. D. Cowan, G. Tesauro, and J. Alspector (editors), *Advances in Neural Information Processing Systems 6*, 51-58, San Mateo, CA: Morgan Kaufmann.

Forrest, S. and Mitchell, M. (1993). Relative building-block fitness and the building-block hypothesis. In D. Whitley (editor), *Foundations of Genetic Algorithms 2*, San Mateo, CA: Morgan Kaufmann.

Mitchell, M., Forrest, S., and Holland, J. H. (1992). The royal road for genetic algorithms: Fitness landscapes and GA performance. In F. J. Varela and P. Bourguine (editors), *Proceedings of the First European Conference on Artificial Life*. Cambridge, MA: MIT Press.

Forrest, S. and Mitchell, M. (1991). The performance of genetic algorithms on Walsh polynomials: Some anomalous results and their explanation. In R. Belew and L. Booker (editors), *Proceedings of the Fourth International Conference on Genetic Algorithms*. San Mateo, CA: Morgan Kaufmann.

Mitchell, M. and Hofstadter, D. R. (1990). The right concept at the right time: How concepts emerge as relevant in response to context-dependent pressures. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Mitchell, M. and Hofstadter, D. R. (1989). The role of computational temperature in a computer model of concepts and analogy-making. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hofstadter, D. R. and Mitchell, M. (1988). *Conceptual slippage and analogy-making: A report on the Copycat project*. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hofstadter, D. R. and Mitchell, M. (1988). Concepts, analogies, and creativity. In *Proceedings of the Canadian Society for Computational Studies of Intelligence*. Edmonton: University of Alberta.

Essays and Book Reviews

Mitchell, M. (2025). Forward to S. Risi, Y. Tang, D. Ha, and R. Miikkulainen, *Neuroevolution: Harnessing Creativity in AI Agent Design*. MIT Press.

Mitchell, M. (2025). Why AI chatbots lie to us. *Science*, July 2025.

Mitchell, M. (2025). Review of *These Strange New Minds* by Christopher Summerfield. *Times Literary Supplement*, May, 2025.

Mitchell, M. (2025). Artificial intelligence learns to reason. *Science*, March 2025.

Mitchell, M. (2024). The metaphors of artificial intelligence. *Science*, November 2024.

Mitchell, M. (2024). The Turing test and our shifting conceptions of intelligence. *Science*, August 2024.

Mitchell, M. (2024). Debates on the nature of artificial general intelligence. *Science*, March 2024.

- Mitchell, M. (2023). AI's challenge of understanding the world. *Science*, November 2023.
- Mitchell, M. (2023). How do we know how smart AI systems are? *Science*, July 2023.
- Mitchell, M. (2022). What does it mean to align AI with human values? *Quanta*, December 2022.
- Mitchell, M. (2021). What does it mean for AI to understand? *Quanta*, December, 2021.
- Jenkins, O. C., Lopresti, D., and Mitchell, M. (2020). Next wave artificial intelligence: Robust, explainable, adaptable, ethical, and accountable. *CCC Quadrennial Papers*, November 2020.
- Flack, J. and Mitchell, M. (2020). Uncertain times. *Aeon*, August, 2020.
- Mitchell, M. (2020). Can GPT-3 make analogies? *Medium*, August, 2020.
- Mitchell, M. (2020). How the analogies we live by shape our thoughts, *Transmissions*, Santa Fe Institute, April, 2020.
- Mitchell, M. (2019). Can a computer ever learn to talk? *OneZero*, November, 2019.
- Mitchell, M. (2019). We shouldn't be scared by 'superintelligent A.I.', *New York Times*, November, 2019.
- Mitchell, M. (2019). Blade Runner is set in November 2019, but what does it say about our future? *The Big Issue*, November, 2019.
- Mitchell, M. (2019). AI can pass standardized tests—but it would fail preschool. *Wired*, September 10, 2019.
- Mitchell, M. (2019). How do you teach a car that a snowman won't walk across the road? And other conundrums of artificial intelligence. *Aeon*, May 31, 2019.
- Mitchell, M. (2018). Artificial intelligence hits the barrier of meaning. *New York Times*, November 5, 2018.
- Mitchell, M. (2014). Adaptive computation: Information, adaptation, and evolution in silico. *Santa Fe Institute Bulletin*, 28(2).
- Mitchell, M. (2014). How can the study of complexity transform our understanding of the world? *Big Questions Online* (<http://bigquestionsonline.com>).
- Mitchell, M. (2008). Visual understanding. *Santa Fe Institute Bulletin*, Spring 2008, pp. 50–54.
- Mitchell, M. (2003). Review of “Conceptual Coordination: How the Mind Orders Experience in Time” by William J. Clancey. *Contemporary Psychology*, 48 (3).
- Mitchell, M. (2002). Review of “A New Kind of Science” by Stephen Wolfram. *Science*, 298, 65–68.
- Mitchell, M. (1997). Review of “Figments of Reality” by Ian Stewart and Jack Cohen. *New Scientist*, August 11, 1997.

Mitchell, M. (1998). Review of “Handbook of Genetic Algorithms” by Lawrence Davis. *Artificial Intelligence*, 100 (1-2), 325-330.

Mitchell, M. (1997). Review of “Darwin’s Dangerous Idea” by Daniel Dennett. *Complexity*, 2 (1), 32–26.

Mitchell, M. (1995). Review of “Out of Control: The Rise of Neo-Biological Civilization” by Kevin Kelly. *Technology Review*, October, 1995.

Mitchell, M. (1993). Computer models of adaptive complex systems. *New Scientist*, February 13, 1993.

Mitchell, M. (1991). Review of “The Dreams of Reason: The Computer and the Rise of the Sciences of Complexity” by Heinz Pagels. In *Bulletin of the Santa Fe Institute*, 6 (1).

Mitchell, M. (1985). Artificial intelligence and the popular press. *Popular Computing*, January, 1985.

Unpublished Preprints

Goldblum, M., Anandkumar, A., Baraniuk, R., Goldstein, T., Cho, K., Lipton, Z.C., Mitchell, M., Nakkiran, P., Welling, M. and Wilson, A.G. (2023). Perspectives on the state and future of deep learning—2023. *arXiv:2312.09323*.

INVITED LECTURES

See <https://melaniemitchell.me/> (“Speaking” tab) for recent invited lectures.

COURSES TAUGHT

Introduction to Complexity (Massive Open Online Course, Santa Fe Institute, 2013-2014, over 25,000 enrolled to date)

Artificial Intelligence (undergraduate and graduate level, University of Michigan, 1991; Portland State University, 2006–present)

Cognitive Science (undergraduate level, University of Michigan, 1991)

Emergent Computation (graduate level, University of New Mexico, 1995)

Evolutionary Computation (undergraduate and graduate level, Oregon Graduate Institute, 2002)

Exploring Complexity in Science and Technology (undergraduate level, Portland State University, 2009–2015)

Machine Learning (undergraduate and graduate level, Oregon Graduate Institute and Portland State University, 2003–2019)

Advanced Topics in Machine Learning (undergraduate and graduate level, Oregon Graduate Institute and Portland State University, 2005, 2009, 2010)

Machine Learning Seminar (graduate level, Portland State University, 2014–2019)

Nonstandard Computation (undergraduate and graduate level, Oregon Graduate Institute and Portland State University, 2004–2005, 2008)

CONFERENCES, WORKSHOPS, AND SUMMER SCHOOLS ORGANIZED

Co-Organizer (with John Krakauer), SFI Working Group on the Nature of Intelligence. August 2024, March 2025, and March 2026.

Co-Organizer (with Jacob Foster), SFI Working Group on Building Diverse Intelligences Through Compositionality and Mechanism Design. March 2026.

Co-Organizer (with Tyler Millhouse, and Melanie Moses), SFI Workshop on AI and the Barrier of Meaning 2. April 2023.

Co-Organizer (with Tyler Millhouse, and Melanie Moses), SFI Workshop on Embodied, Situated, and Grounded Intelligence: Implications for AI. April 2022.

Co-Organizer (with Tyler Millhouse, and Melanie Moses), SFI Workshop on Collective Intelligence in Natural and Artificial Systems. August, 2021.

Co-Organizer (with Stephanie Forrest, Risto Miikula, Tyler Millhouse, and Melanie Moses), SFI Workshop on Frontiers of Evolutionary Computing. July 2021

Co-Organizer (with Yuhuai Wu, Kshitij Bansal, Wenda Li, David McAllester, and John Harrison), ICLR Workshop on The Role of Mathematical Reasoning in General Artificial Intelligence. May 2021.

Co-Organizer (with Melanie Moses), SFI Workshop on Foundations of Intelligence in Natural and Artificial Systems, March 2021.

Co-Organizer (with François Chollet and Kevin Ellis), AAAI Fall Symposium on Conceptual Abstraction and Analogy in Natural and Artificial Intelligence. November 2020.

Co-Organizer (with Barbara Grosz and Dawn Song), workshop on Artificial Intelligence and the Barrier of Meaning. Santa Fe Institute, Santa Fe, NM, October 2018.

Organizer, short course on Exploring Complexity in Science and Technology from a Santa Fe Institute Perspective. Portland, OR, 2010; Albuquerque, NM, May, 2011; Stanford, CA, Sept. 2012; Austin, TX, Sept. 2013; Santa Fe NM, 2015-2016.

Co-Organizer (with SFI and the Krasnow Institute), short course on The Science of Complexity: Understanding the Financial Crisis, Washington, DC, May, 2012.

Co-Organizer (with Jo Ann Wise and Ralph Greenspan), workshop on Shared Organizing Principles in the Computing and Biological Sciences. National Science Foundation, Washington, DC. 2010.

Co-Organizer (with Garrett Kenyon, Ilya Nemenman, and Chris Wood), workshop on Principles of Biological Computation. Santa Fe Institute, Santa Fe, NM. 2008.

Co-Organizer (with Garrett Kenyon), workshop on High-Level Perception and Low-Level Vision: Bridging the Semantic Gap. Santa Fe Institute, Santa Fe, NM. 2007.

Director, Complex Systems Summer School, Santa Fe Institute. 1999-2002; 2004–2005 (1999, Co-director with Lynn Stein; 2000, Co-director with Ray Goldstein).

Co-Director (with Imre Kondor), Central European Complex Systems Summer School, Central European University, Budapest. 2000.

Co-Organizer (with Lashon Booker, Stephanie Forrest, and Rick Riolo), Holland Fest meeting (conference in honor of John Holland's 70th birthday). University of Michigan, Ann Arbor, MI. 1999.

Co-Organizer (with Randall Beer), workshop on Dynamics, Computation, and Cognition. Santa Fe Institute. May 1996.

Co-Organizer (with Terry Jones and Una-May O'Reilly), workshop on Biological and Computational Landscapes, Santa Fe Institute. July 1995.

Organizer, working group on Theoretical Foundations of Genetic Algorithms, Santa Fe Institute. January 11–13, 1994, May-June, 1994.

Co-Organizer (with Richard Belew), workshop on Behavioral Plasticity in Evolving Populations: Models and Algorithms, Santa Fe Institute. July 1993.

Organizer, workshop on Learning and Adaptation in Robots and Situated Agents, Santa Fe Institute. May 1993.

Co-Organizer (with Nils Nilsson), workshop on Reinforcement Learning in Robotics, Santa Fe Institute. March 1993.

Organizer, workshop on Computation, Dynamical Systems, and Learning, Santa Fe Institute. November 1992.

PUBLIC OUTREACH

Melanie Mitchell: *AI: A Guide for Thinking Humans*, Substack Newsletter (aiguide.substack.com), over 58,000 subscribers, 2023–present.

Melanie Mitchell and Abha Eli Phoboo, co-hosts: *The Nature of Intelligence*. Six-episode series on the Complexity Podcast from the Santa Fe Institute, 2024.

Ida Momennejad, John Krakauer, and Melanie Mitchell, co-hosts: *The Learning Salon*. Online salon focused on human and machine learning (<https://www.youtube.com/@learningsalon3940>)

POSTDOCS AND STUDENTS SUPERVISED

Postdocs:

Shuhao Fu, Santa Fe Institute, 20205
Arseny Moskvichev, Santa Fe Institute, 2022–2024.
Tyler Millhouse, Santa Fe Institute, 2021–2022.
Efsun Sarioglu, Portland State University, 2015–2016
Manuel Marques-Pita, Portland State University, 2008–2009
Ludo Pagie, Santa Fe Institute, 2000–2002

Ph.D. Students

Max Quinn, Ph.D. 2021, Portland State University.
Sheng Lundquist, Ph.D. 2020, Portland State University
Anthony Rhodes, Ph.D. 2020, Portland State University
Will Landecker, Ph.D. 2014, Portland State University
Ralf Jüngling, Ph.D., 2013, Portland State University
Michael Thomure, Ph.D., 2013, Portland State University
Martin Cenek, Ph.D., 2011, Portland State University
Payel Ghosh, Ph.D., 2010, Portland State University
Wim Hordijk, Ph.D., 1999, University of New Mexico/Santa Fe Institute
(co-advisors James Crutchfield and Stephanie Forrest)
Rajarshi Das, Ph.D., 1996, Colorado State University/Santa Fe Institute
(co-advisors James Crutchfield and Darrell Whitley)

M.S. Theses

Sharad Kumar, M.S., 2018, Portland State University
Andrew Cleland, M.S., 2018, Portland State University
Erik Conser, M.S., 2017, Portland State University
Naomi Dickerson, M.S., 2017, Portland State University
Lewis Coates, M.S., 2016, Portland State University
Kendall Stewart, M.S. 2015, Portland State University
Clint Olson, M.S. 2015, Portland State University
Joanna Solman, M.S., 2014, Portland State University
George Dittmar, M.S., 2013, Portland State University
Karan Sharma, M.S., 2012, Portland State University
Dan Coates, M.S., 2009, Portland State University
Lanfranco Muzi, M.S., 2009, Portland State University
Davis Stevenson, M.S. 2007, Portland State University
Nathan Williams, M.S., 2004, Oregon Health & Science University

Undergraduate and Postbac Interns:

Ryan Yi, Santa Fe Institute, 2025
Claas Beger, Santa Fe Institute, 2025
Yiyang Liu, Santa Fe Institute, 2025
Alessandro Palmarini, Santa Fe Institute, 2023-2024
Ky-Vinh Mai, UC Irvine, 2023
Andrew Geyko, University of New Mexico, 2023
Yutaro Shimizu, Minerva University, 2022
Daniel Cotayo, Florida International University, 2022
Victor Odouard, Santa Fe Institute, 2021-2022
Julia Beckwith, Scripps College, 2020
Kennedy Hahn, Portland State University, 2018-2019
Chandler Watson, Stanford University, 2018
Evan Roche, Lewis & Clark College, 2016
Rory Soiffer, University of Washington, 2016
Favian Rahman, Carnegie Mellon University, 2014
Max Boddy, Reed College, 2013
Eben Wood, Portland State University, 2012-2013
Max Orhai, Portland State University, 2012-2013
Jennifer Meneghin, Portland State University, 2006
Michael Thomure, Portland State University, 2004–2005
Jonathan Carlson, Dartmouth College, 2003
Justin Werfel, Princeton University, 1998
Alex Wo, Harvard University, 1997
Elizabeth Ayer, Duke University, 1995
Adam Messinger, Willamette University, 1994
Peter Hrabec, Santa Fe Institute, 1993

High School Interns:

Anna Alice Patelli, Albuquerque Academy, Albuquerque, NM 2026
Shrey Poshia, Santa Fe Prep, Santa Fe, NM, 2021
Joaquin Bas, Santa Fe Prep, Santa Fe, NM, 2020
Robin Tan, Jesuit High School, Portland, OR, 2016
Bryan Lee, Westview High School, Beaverton, OR, 2016
Chandler Watson, Oregon Episcopal School, Portland, OR 2014-2015
Vicki Niu, Lincoln High School, Portland, OR, 2013
Preetha Velu, Jesuit High School, Portland, OR, 2013

RECENT PROFESSIONAL SERVICE (since 2020)

Co-Chair, Science Board, Santa Fe Institute (2019–2025)
Member of Science Steering Committee, Santa Fe Institute (2020–present)
Member, Open Strategic Partnership Board, OpenEuroLLM Consortium (2025–present)
Member, Science and Security Board, Bulletin of the Atomic Scientists (2025–present)
Member, Advisory Board, Scripps Research Translational Institute (2023–2024)
Member, Advisory Board, *AI Magazine*, 2020–present.

Member, Editorial Board, *Collective Intelligence*, 2020–present.

Member, AI100 Study Panel, 2020–2021.

Member, Computing Community Consortium (CCC) Council, 2018–2021.

Member of Program Committee for 2020 AAAI Conference

RECENT K–12 EDUCATIONAL ACTIVITIES (since 2020)

Advisor, Everyday AI program (MIT), 2021–present.

Mentor, Institute for Computing in Research, Santa Fe, NM, 2020–2023.