



INTELLIGENCE, NATURAL AND ARTIFICIAL CITATIONS
June 11–12, 2018 • Brooklyn, New York

prepared by Nancy Kleinrock nk33@cornell.edu

Citations (books, films, and TV series, listed by speaker)

Len Kleinrock		
Paine, Chris (director)	<i>Do You Trust This Computer?</i>	2018
Noam Brown		
Campbell, Martin (director)	<i>Casino Royale</i>	2006
Dahl, John (director)	<i>Rounders</i>	1998
Anthony Zador		
Heinlein, Robert	<i>Methuselah's Children</i>	1958
Minsky, Marvin, and Seymour Papert	<i>Perceptrons: An Introduction to Computational Geometry</i>	1969
Cameron, James (director)	<i>The Terminator</i>	1984
Jeff Jonas		
Mezrich, Ben	<i>Bringing Down the House</i>	2003
Luketic, Robert (director)	<i>21</i>	2008
Stanley, Kenneth		
Stanley, Kenneth, and Joel Lehman	<i>Why Greatness Cannot Be Planned: The Myth of the Objective</i>	2015
Nolfi, Stefano, and Dario Floreano	<i>Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines</i>	2000
Gary Marcus		
Kubrick, Stanley (director, producer, cowriter)	<i>2001: A Space Odyssey</i>	1968
Jonze, Spike (director)	<i>Her</i>	2013
Marcus, Gary	<i>The Algebraic Mind: Integrating Connectionism and Cognitive Science</i>	2001
Sacks, Oliver	<i>The Man Who Mistook His Wife for a Hat: And Other Clinical Tales</i>	2006
Pearl, Judea	<i>The Book of Why: The New Science of Cause and Effect</i>	2018
Pinker, Steven	<i>Language Learnability and Language Development</i>	1984
Fodor, Jerry	<i>The Language of Thought</i>	1975
Doug Lenat		
Kahneman, Daniel	<i>Thinking, Fast and Slow</i>	2011
Ben Kuipers		
Schreier, Jake (director)	<i>Robot & Frank</i>	2012
Fukuyama, Francis	<i>Trust: The Social Virtues and the Creation of Prosperity</i>	1996
Russell, Stuart, and Peter Norvig	<i>Artificial Intelligence: A Modern Approach, 3rd</i>	2010

	<i>Edition</i>	
Singer, Peter	<i>The Expanding Circle: Ethics and Sociology</i>	1981
Wright, Robert	<i>Nonzero: The Logic of Human Destiny</i>	1999
Pinker, Steven	<i>The Better Angels of Our Nature: Why Violence Has Declined</i>	2011
Pinker, Steven	<i>Enlightenment Now: The Case for Reason, Science, Humanism, and Progress</i>	2018
Beauchamp, Tom, and James Childress	<i>Principles of Biomedical Ethics, 6th Edition</i>	2008
Winkler, Adam	<i>We the Corporations: How American Businesses Won Their Civil Rights</i>	2018
Cameron, James (director)	<i>Terminator 2: Judgment Day</i>	1991
Elias Bareinboim		
Pearl, Judea	<i>The Book of Why: The New Science of Cause and Effect</i>	2018
Pearl, Judea	<i>Causality: Models, Reasoning, and Inference, 2nd Edition</i>	2009
Halpern, Joseph	<i>Actual Causality</i>	2016
Fisher, Ronald	<i>The Design of Experiments</i>	1935
Guy Hoffman		
Schreier, Jake (director)	<i>Robot & Frank</i>	2012
Stebbins, Genevieve	<i>Delsarte System of Dramatic Expression</i>	1886
Darwin, Charles	<i>The Expression of the Emotions in Man and Animals</i>	1872

Online Resources (by speaker, in order of appearance)

Noam Brown

- <https://www.cs.cmu.edu/~noamb/research.html> — List of Brown's academic publications.
- <https://www.cs.cmu.edu/~sandholm/RIHoldEm.ISD.aaai05proceedings.pdf> — Full text of Gilpin, A., and T. Sandholm. 2005. Optimal Rhode Island hold'em poker. *AAAI'05 Proceedings of the 20th Annual National Conference on Artificial Intelligence* 4:1685–1685.
- <http://martin.zinkevich.org/publications/regretpoker.pdf> — Full text of Zinkevich, M., et al. 2007. Regret minimization in games with incomplete information. *Advances in Neural Information Processing Systems* 20:905–912.
- <https://www.cs.rutgers.edu/~mlittman/papers/cg00-poker.pdf> — Full text of Shi, J., and M.L. Littman. 2001. Abstraction methods for game theoretic poker. Pages 333–345 in *Computers and Games*. Springer-Verlag.
- <http://poker.cs.ualberta.ca/publications/IJCAI03.pdf> — Full text of Billings, D., et al. 2003. Approximating game-theoretic optimal strategies for full-scale poker. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*.
- <https://www.itu.dk/~trbj/papers/tartanian.pdf> — Full text of Gilpin, A., et al. 2008. A heads-up no-limit Texas hold'em poker player: Discretized betting models and automatically generated equilibrium-finding programs. In *Proceeding so the 9th International Conference on Autonomous Agents and Multiagent Systems*.
- <https://webdocs.cs.ualberta.ca/~holte/Publications/AAAI-2012-poker.pdf> — Full text of Hawkin, J., et al. 2012. Using windows to generate action abstractions in extensive-form games. In *AAAI Conference on Artificial Intelligence*.
- https://www.cs.cmu.edu/~sandholm/regret_transfer.aaai14.pdf — Full text of Brown, N., and T. Sandholm. 2014. Regret transfer and parameter optimization. In *AAAI Conference on Artificial Intelligence*.
- <https://ijcai.org/Proceedings/09/Papers/055.pdf> — Full text of Schnizlein, D., et al. 2009. Probabilistic state translation in extensive games with large action sets. In *IJCAI*.
- <https://www.cs.cmu.edu/~sandholm/reverse%20mapping.ijcai13.pdf> — Full text of Ganzfried, S., and T. Sandholm. 2013. Action translation in extensive-form games with large action spaces: Axioms,

- paradoxes, and the pseudo-harmonic mapping. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.295.2143&rep=rep1&type=pdf> — Full text of Johanson, M., et al. 2013. Evaluating state-space abstractions in extensive-form games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*.
 - https://www.cs.cmu.edu/~sandholm/potential-aware_imperfect-recall_aai14.pdf — Full text of Ganzfried, S., and T. Sandholm. 2014. Potential-aware imperfect-recall abstraction with earth mover's distance in imperfect-information games. In *AAAI Conference on Artificial Intelligence*.
 - <https://arxiv.org/pdf/1205.0622.pdf> — Full arXiv version of Lanctot, M., et al. 2012. No-regret learning in extensive-form games with imperfect recall. In *International Conference on Machine Learning*.
 - <http://poker.cs.ualberta.ca/publications/aaai2014-cfrd.pdf> — Full text of Burch, N., et al. 2014. Solving imperfect games using decomposition. In *AAAI Conference on Artificial Intelligence*.
 - <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12102/11636> — Full text of Moravcik, M., et al. 2016. Refining subgames in large imperfect information games. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*.
 - <https://www.cs.cmu.edu/~noamb/papers/16-IJCAI-Prune.pdf> — Full text of Brown, N., et al. 2017. Dynamic thresholding and pruning for regret minimization. In *AAAI Conference on Artificial Intelligence*.
 - https://www.cs.cmu.edu/~sandholm/endgame_aamas15_fromACM.pdf — Full text of Ganzfried, S., and T. Sandholm. Endgame solving in large imperfect-information games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*.
 - <https://www.cs.cmu.edu/~noamb/papers/17-NIPS-Safe.pdf> — Full text of Brown, N., and T. Sandholm. 2017. Safe and nested subgame solving for imperfect-information games. In *Proceedings of the Conference on Neural Information Processing Systems*.
 - <https://arxiv.org/pdf/1805.08195.pdf> — Full preprint of Brown's most recent publication: Brown, N., et al. 2018. Depth-limited solving for imperfect-information games. *arXiv:1805.08195*.
 - <https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now> — Article by Andrew Ng, first appearing in *Harvard Business Review* on November 9, 2016, entitled "What artificial intelligence can and can't do right now."

Bo Zhu

- http://mi.eng.cam.ac.uk/~kmk/presentations/TutorialIC_Sep2015_part2_Knill.pdf — Tutorial slides by Kate Knill, first appearing in September 2015, entitled "(Deep) neural networks for speech processing."
- <https://www.nature.com/articles/nn2077> — Abstract and access options for Doya, K. 2008. Modulators of decision making. *Nature Neuroscience* **11**:410–416.
- <https://www.nature.com/articles/nn2077> — Abstract and access options for Volodymyr, M., et al. 2015. Human-level control through deep reinforcement. *Nature* **518**:529–533.
- <https://arxiv.org/pdf/1704.08841.pdf> — Full arXiv version of Zhu, B., et al. 2018. Image reconstruction by domain-transform manifold learning. *Nature* **555**:487–492.
- <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2864603/> — Full text of Sasaki, Y., et al. 2010. Advances in perceptual learning and plasticity. *Nature Reviews Neuroscience* **11**:53–60.
- <http://vision.ustc.edu.cn/papers/Visual%20perceptual%20learning.pdf> — Full text of Lu, Z.-L., et al. 2011. Visual perception learning. *Neurobiology of Learning and Memory* **95**:145–151.
- https://www.cs.cmu.edu/~epxing/Class/10715/reading/Kornick_et_al.pdf — Full text of Hornik, K. 1989. Multilayer feedforward networks are universal approximations. *Neural Networks* **2**:359–366.
- https://web.eecs.umich.edu/~cscott/smlrg/approx_by_superposition.pdf — Full text of Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems* **2**:303–314.
- <http://cds.ismrm.org/ismrm-2001/PDF1/0004.pdf> — Full text of Pruessmann, K.P., et al. 2001. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance in Medicine* **46**:638–651.
- <https://hal.archives-ouvertes.fr/hal-01813870/document> — Full text of Guerquin-Kern, M., et al. 2011. A fast wavelet-based reconstruction method for magnetic resonance imaging. *IEEE Transactions on Medical Imaging* **30**(9):1649–1660.
- <https://www.sciencedirect.com/science/article/pii/S0377042711005188> — Full text of Hansen, P.C., and M. Saxild-Hansen. 2012. AIR tools: A MATLAB package of algebraic iterative reconstruction techniques. *Journal of Computational and Applied Mathematics* **236**:2167–2178.

- <http://index.miramsmart.com/ismrm2015/PDFfiles/2486.pdf> — Full text of Uecker, M., et al. Berkeley advanced reconstruction toolbox. *Proceedings of the International Meeting of the Society of Magnetic Resonance in Medicine* **23**:2486.

Anthony Zador

- <http://zadorlab.cshl.edu/> — List of Zador's publications.
- <http://www.cse.chalmers.se/~coquand/AUTOMATA/mcp.pdf> — Full text of McCulloch, W.S., and W.H. Pitts. 1943. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* **5**(4):115–133.
- <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.335.3398&rep=rep1&type=pdf> — Full text of Rosenblatt, F. 1958. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review* **65**:386–408.
- <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1363130> — Full text of Hubel, D.H., and T.N. Wiesel. 1959. Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology* **148**:574–591.
- <http://www.cs.toronto.edu/~hinton/absps/naturebp.pdf> — Full text of Rumelhart, D.E., et al. 1986. Learning representations by back-propagating errors. *Nature* **323**:533–536.
- <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1359523/> — Full text of Hubel, D.H., and T.N. Wiesel. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology* **160**:106–154.
- <https://academic.oup.com/nar/article/46/4/e22/4668654> — Full text of Zador's paper: Chen, X., et al. 2018. Efficient *in situ* barcode sequencing using padlock probe-based BaristaSeq. *Nucleic Acids Research* **46**:e22.

Jeff Jonas

- <https://senzing.com/wp-content/uploads/Senzing-GDPR-Missing-Link-Report-FINAL-FINAL.pdf> — Publication by Senzing, entitled “Finding the missing link in GDPR compliance.”
- <http://www.nbcnews.com/id/6192603/ns/health-arthritis/t/report-vioxx-linked-thousands-deaths> — Article first appearing on NBC News on October 6, 2004, entitled “Report: Vioxx linked to thousands of deaths.”

David Gunning

- <https://www.darpa.mil/program/explainable-artificial-intelligence> — Learn more about DARPA's Explainable Artificial Intelligence (XAI) program.
- <https://www.nvidia.cn/content/tesla/pdf/machine-learning/imagenet-classification-with-deep-convolutional-nn.pdf> — Full text of Krizhevsky, A., et al. 2012. ImageNet classification with deep convolutional neural networks. *Proceedings of Advances in Neural Information Processing Systems* **25**:1090–1098.
- <https://arxiv.org/pdf/1802.08129.pdf> — Full text of Park, D.H., et al. 2018. Multimodal explanations: Justifying decisions and pointing to the evidence. *arXiv*:1802.08129.
- <https://arxiv.org/pdf/1703.10631.pdf> — Full text of Kim, J., et al. 2017. Show, attend, control, and justify: Interpretable learning for self-driving cars. *arXiv*:1703.10631.
- https://people.eecs.berkeley.edu/~jfc/papers/17/iccv-final-interpretible_learning_for_self_driving_cars.pdf — Full text of Kim, J., and J. Canny. 2018. Interpretable learning for self-driving cars by visualizing causal attention. In *Proceedings of the IEEE International Conference on Computer Vision*.

Kenneth Stanley

- <https://www.cs.ucf.edu/~kstanley> — A synopsis of what is going on in the Stanley research group.
- <https://link.springer.com/article/10.1007/BF01411376> — Full text of Radcliffe, N.J. 1993. Genetic set recombination and its application to neural network topology optimisation. *Neural Computing and Applications* **1**:67–90.
- http://axon.cs.byu.edu/~dan/778/papers/NeuroEvolution/stanley3**.pdf — Full text of Stanley, K.O., et al. A hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life* **15**:185–212.
- <http://jeffclune.com/publications/2011-CluneEtAl-IndirectEncodingAcrossRegularityContinuum-IEEE-TEC.pdf> — Full text of Stanley's paper: Clune, J., et al. 2011. On the performance of indirect encoding across the continuum of regularity. In *IEEE Transactions of on Evolutionary Computation*.

- http://eplex.cs.ucf.edu/papers/verbancsics_jmlr10a.pdf — Full text of Verbancsics, P., and K.O. Stanley. 2010. Evolving static representations for task transfer. *Journal of Machine Learning Research* **11**:1737–1769.
- http://eplex.cs.ucf.edu/papers/gauci_nc10.pdf — Full text of Gauci, J., and K.O. Stanley. 2010. Autonomous evolution of topographic regularities in artificial neural networks. *Neural Computation* **22**(7):1860–1898.
- http://eplex.cs.ucf.edu/papers/lehman_ecj11.pdf — Full text of Lehman, J., and K.O. Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation* **19**:189–223.
- http://eplex.cs.ucf.edu/papers/lehman_gecco11.pdf — Full text of Lehman, J., and K.O. Stanley. 2011. Evolving a diversity of creatures through novelty search and local competition. In *Proceedings of the Genetic and Evolutionary Computation Conference*.
- <https://arxiv.org/pdf/1407.3501.pdf> — Full arXiv version of Cully, A., et al. 2015. Robots that can adapt like animals. *Nature* **521**:503–507.
- <https://arxiv.org/pdf/1703.03864.pdf> — Full text of Salimans, T., et al. 2017. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv*:1703.03864.
- <https://arxiv.org/pdf/1712.06567.pdf> — Full text of Stanley's paper: Such, F.P., et al. 2017. Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arXiv*:1712.06567.
- <https://arxiv.org/pdf/1712.06560.pdf> — Full text of Stanley's paper: Conti, E., et al. 2018. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. *arXiv*:1712.06560v2.
- <https://arxiv.org/pdf/1802.01548.pdf> — Full text of Real, E., et al. 2018. Regularized evolution for image classifier architecture search. *arXiv*:1802.01548.
- <https://arxiv.org/pdf/1803.03745.pdf> — Full text of Liang, J., et al. 2018. Evolutionary architecture search for deep multitask networks. *arXiv*:1803.03745.
- <https://www.sciencedirect.com/science/article/pii/S0893608000000320> — Abstract and access options for Floreano, D., and J. Urzelai. 2000. Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks* **13**:431–443.
- http://eplex.cs.ucf.edu/papers/risi_sab10.pdf — Full text of Risi, S., and K.O. Stanley. 2010. Indirectly encoding neural plasticity as a pattern of local rules. In *Proceedings of the 11th International Conference on Simulation of Adaptive Behavior*.
- http://eplex.cs.ucf.edu/papers/risi_gecco11.pdf — Full text of Risi, S., and K.O. Stanley. 2011. Enhancing ES-HyperNEAT to evolve more complex regular neural networks. In *Proceedings of the Genetic and Evolutionary Computation Conference*.
- <https://eng.uber.com/deep-neuroevolution> — Article by Stanley and Jeff Clune, first appearing in Uber Engineering Updates on December 18, 2017, entitled “Welcoming the era of deep neuroevolution.”

Gary Marcus

- http://www.psych.nyu.edu/gary/marcus_pubs.html — Listing of Marcus's books, academic publications, articles, chapters, reviews, essays, and op eds.
- <http://www.newsweek.com/robots-can-now-read-better-humans-putting-millions-jobs-risk-781393> — Article by Anthony Cuthbertson, first appearing in *Newsweek* on January 15, 2018, entitled “Robots can now read better than humans, putting millions of jobs at risk.”
- <https://www.recode.net/2018/1/19/16911180/sundar-pichai-google-fire-electricity-ai> — Article by Theodore Schleifer, first appearing on Recode on January 19, 2018, entitled “Google CEO Sundar Pichai says AI is more profound than electricity and fire.”
- <https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124> — Article by Henry Kissinger, first appearing in *The Atlantic* in June 2018, entitled “How the Enlightenment ends: Philosophically, intellectually—in every way—human society is unprepared for the rise of artificial intelligence.”
- <https://www.technologyreview.com/s/609141/alphago-zero-shows-machines-can-become-superhuman-without-any-help> — Article by Will Knight, first appearing on Technology Review on October 18, 2017, entitled “AlphaGo Zero shows machines can become superhuman without any help.”
- <https://www.nytimes.com/2018/05/18/opinion/artificial-intelligence-challenges.html> — Article by Gary Marcus and Ernest Davis, first appearing in the *New York Times* on May 18, 2018, entitled “A.I. is harder than you think.”

- <https://arxiv.org/pdf/1801.00631.pdf> — Full text of Marcus, G. 2018. Deep learning: A critical appraisal. *arXiv*:1801.00631.
- <https://www.wired.com/story/greedy-brittle-opaque-and-shallow-the-downsides-to-deep-learning> — Article by Jason Pontin, first appearing in *Wired* on February 2, 2018, entitled “Greedy, brittle, opaque, and shallow: The downsides to deep learning.”
- <https://arxiv.org/pdf/1707.07397.pdf> — Full text of Athalye, A., et al. 2017. Synthesizing robust adversarial examples. *arXiv*:1707.07397.
- <https://arxiv.org/pdf/1712.09665.pdf> — Full text of Brown, T., et al. 2017. Adversarial patch. *arXiv*:1712.09665.
- <https://cacm.acm.org/magazines/2015/9/191169-commonsense-reasoning-and-commonsense-knowledge-in-artificial-intelligence/abstract> — Abstract and access options for Marcus, G., and E. Davis. 2015. Commonsense reasoning and commonsense artificial intelligence. *Communications of the ACM* **58**(9):92–103.
- <https://www.theguardian.com/technology/2016/aug/15/roomba-robot-vacuum-poopocalypse-facebook-post> — Article by Olivia Solon, first appearing in the *Guardian* on August 15, 2016, entitled “Roomba creator responds to reports of ‘poopocalypse’: ‘We see this a lot.’”
- <https://www.newyorker.com/news/news-desk/is-deep-learning-a-revolution-in-artificial-intelligence> — Article by Gary Marcus, first appearing in *The New Yorker* on November 15, 2012, entitled “Is ‘deep learning’ a revolution in artificial intelligence?”
- <http://www.psych.nyu.edu/gary/marcusArticles/marcus%20et%20al%201999%20science.pdf> — Full text of Marcus, G.F., et al. 1999. Rule learning by seven-month-old infants. *Science* **283**:77–80.
- <https://pdfs.semanticscholar.org/accc/59256e239a4d297a98195d8ed060b74dce26.pdf> — Full text of Spelke, E.S., et al. 1994. Early knowledge of object motion: Continuity and inertia. *Cognition* **51**:131–176.
- <https://www.youtube.com/watch?v=g0TaYhjpOfo> — Video, first posted by *IEEE Spectrum* on June 6, 2015, entitled “A compilation of robots falling down at the DARPA Robotics Challenge.”
- <https://medium.com/intuitionmachine/how-to-grow-the-innate-machinery-for-agi-a5cbbd755eae> — Article by Carlos Perez, first appearing on Medium on November 4, 2017, entitled “How to grow the innate machinery for AGI.”
- <https://arxiv.org/pdf/1801.05667.pdf> — Full text of Marcus, G. 2018. Innateness, AlphaZero, and artificial intelligence. *arXiv*:1801.05667.

Doug Lenat

- <https://arxiv.org/pdf/1801.00631.pdf> — Full text of Marcus, G. 2018. Deep learning: A critical appraisal. *arXiv*:1801.00631.
- <http://time.com/5304611/france-fake-news-law-macron> — Article by Ciara Nugent, first appearing in *Time* on June 7, 2018, entitled “France is voting on a law banning fake news. Here’s how it could work.”
- <https://www.theguardian.com/world/2018/jun/07/france-macron-fake-news-law-criticised-parliament> — Article by Angelique Chrisafis, first appearing in the *Guardian* on June 8, 2018, entitled “French MPs criticize ‘hasty and ineffective’ fake news law.”
- [https://faculty.washington.edu/jdb/345/345%20Articles/lyengar%20%26%20Lepper%20\(2000\).pdf](https://faculty.washington.edu/jdb/345/345%20Articles/lyengar%20%26%20Lepper%20(2000).pdf) — Full text of lyengar, S.S., and M.R. Lepper. 2000. When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology* **79**(6):995–1006.
- <https://stanford.app.box.com/s/yohfziywajw3nmwxo7d3ammndihibe7g> — Full text of Davidai, S., et al. 2012. The meaning of default options for potential organ donors. *Proceedings of the National Academy of Sciences* **109**(38):15201–15205.
- <http://www.health.thesfile.com/children/vaccinations/statistical-risk-of-suffering-a-severe-case-of-vaccine-preventable-disease/> — Review by The S File of *The Vaccine Book*, by Robert Sears (2007), entitled “Statistical risk of suffering a severe case of vaccine-preventable disease.”
- <https://www.sas.upenn.edu/~baron/papers.htm/vac.html> — Full text of Ritov, I., and J. Baron. 1990. Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making* **3**:263–277.
- <https://pdfs.semanticscholar.org/e456/4b88ca2349962a707b76be4c75076ad6bd43.pdf> — Full text of Arkes, H.R. and C. Blumer. 1985. The psychology of sunk cost. *Organizational Behavior and Human Decision Processes* **35**:124–140.
- <http://web.missouri.edu/~segerti/capstone/choicesvalues.pdf> — Full text of Kahneman, D., and A. Tversky. 1984. Choices, values, and frames. *American Psychologist* **39**:341–350.

- https://en.wikipedia.org/wiki/Simpson%27s_paradox — Explanation of Simpson's paradox (otherwise known as the Yule-Simpson effect).
- <https://www.nature.com/news/how-scientists-fool-themselves-and-how-they-can-stop-1.18517> — Full text of Nuzzo, R. 2015. How scientists fool themselves—and how they can stop. *Nature* **526** (article).
- <https://arxiv.org/pdf/1412.1897.pdf> — Full *arXiv* version of Nguyen, A., et al. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Ben Kuipers

- <https://web.eecs.umich.edu/~kuipers/vita.pdf> — List of Kuipers's books and academic publications.
- <https://www.youtube.com/watch?v=eQxUW4B622E> — Videoclip from the movie *Robot & Frank*.
- <https://www.youtube.com/watch?v=3yXwPfvlt4> — Another videoclip from the movie *Robot & Frank*.
- <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.470.8322&rep=rep1&type=pdf> — Full text of Rousseau, D.M., et al. 1998. Not so different after all: A cross-discipline view of trust. *Academy of Management Review* **23**:393–404.
- <http://lefft.xyz/psycholingAU16/readings/grice1975-logic-and-conversation.pdf> — Full text of Grice, Paul. 1975. Logic and conversation. Pages 41–58 in *Syntax and Semantics, Vol. 3, Speech Acts*, Peter Cole and Jerry L. Morgan, editors.

Rohini Rewari, Robie Samanta Roy, and Vijay Sankaran

- <https://www.hsdl.org/?abstract&did=726163> — Abstract and access link for U.S. Department of Defense Directive 3000.09: Autonomy in Weapons Systems, November 21, 2012, as mentioned by Roy.

Elias Bareinboim

- <https://www.cs.purdue.edu/homes/eb> — Scroll down for a listing of Bareinboim's academic publications.
- <http://thechart.blogs.cnn.com/2013/08/15/study-heavy-coffee-drinking-in-people-under-55-linked-to-early-death> — Article first appearing on CNN's blog The Chart on August 15, 2013, entitled "Study: Heavy coffee drinking in people under 55 linked to early death."
- http://www.infekt.ch/content/uploads/2013/11/jc_juli08_rauch.pdf — Full text of Lopez-Garcia, E., et al. 2008. The relationship of coffee consumption with mortality. *Annals of Internal Medicine* **148**:904–914.
- <https://www.cbsnews.com/news/alcohol-causes-20000-cancer-deaths-in-the-us-annually> — Article by Michelle Castillo, first appearing on *CBS News* on February 15, 2013, entitled "Alcohol causes 20,000 cancer deaths in the U.S. annually."
- <https://www.sciencedaily.com/releases/2008/02/080212174615.htm> — Article by the University Health Network, first appearing in *ScienceDaily* on February 13, 2008, entitled "One drink of red wine or alcohol is relaxing to circulation, but two drinks are stressful."
- <https://content.time.com/time/magazine/article/0,9171,2017200,00.html> — Article by John Cloud, first appearing in *Time* on August 30, 2010, entitled "Why do heavy drinkers outlive nondrinkers?"
- <https://well.blogs.nytimes.com/2012/12/17/grapefruit-is-a-culprit-in-more-drug-reactions> — Article by Roni Caryn Rabin, first appearing in the *New York Times* on December 17, 2012, entitled "Grapefruit is a culprit in more drug reactions."
- <https://www.nejm.org/doi/full/10.1056/nejmoa1307352> — Full text of Bao, Y., et al. 2013. Association of nut consumption with total and cause-specific mortality. *New England Journal of Medicine* **369**:2001–2011.
- <http://www.pnas.org/content/113/27/7345> — Full text of Bareinboim, E., and J. Pearl. 2016. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences* **113**(27):7345–7352.
- <https://www.ijcai.org/proceedings/2017/0186.pdf> — Full text of Zhang, J., and E. Bareinboim. 2017. Transfer learning in multi-armed bandits: A causal approach. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*.
- <https://www.cs.purdue.edu/homes/eb/mdp-causal.pdf> — Full text of Zhang, J., and E. Bareinboim. 2016. Markov decision processes with unobserved confounders: A causal approach.

Erik Andrejko

- <https://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf> — Full text of Mikolov, T., et al. 2013. Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems* 3111–3119.

- <https://gking.harvard.edu/files/gking/files/0314policyforumff.pdf> — Full text of Lazer, D., et al. 2014. The parable of Google flu: Traps in big data analysis. *Science* **343**:1203–1205.
- <https://nlp.stanford.edu/projects/glove> — Learn about the Global Vectors for Word Representation (GloVe) project for unsupervised learning to obtain vector representations for words.
- <https://www.wikidata.org/wiki/Q504133> — Wikidata page for “jalapeño.”
- <https://blog.openai.com/unsupervised-sentiment-neuron> — OpenAI blog post, first appearing on April 6, 2017, entitled “Unsupervised sentiment neuron.”
- <https://newsroom.clevelandclinic.org/2017/04/22/cleveland-clinic-study-finds-obesity-top-cause-preventable-life-years-lost> — News release from the Cleveland Clinic, first appearing on April 22, 2017, entitled “Cleveland Clinic study finds obesity as top cause of preventable life-years lost.”

Guy Hoffman

- <http://guyhoffman.com/category/publications> — Listing of Hoffman’s academic publications.
- https://www.youtube.com/watch?v=okFoKJK_N3w — Video first posted by Hoffman’s Human–Robot COllaboartion and Companionship Lab on August 22, 2017, entitled “Blossom—a handcrafted social robot—soft inside and out.”
- <https://www.economist.com/leaders/2014/03/29/rise-of-the-robots> — Article first appearing in *The Economist* on March 29, 2014, entitled “Rise of the robots: Prepare for a robot invasion. It will change the way people think about technology.”
- <http://www.businessinsider.com/robot-that-can-understand-emotions-2014-6> — Article by Quinton O’Reilly, first appearing in *Business Insider* on June 6, 2014, entitled “Meet the ‘human-like’ robot that can understand people’s emotions.”
- <https://www.cnn.com/2015/06/22/tech/pepper-robot-sold-out/index.html> — Article by Angad Singh, first appearing on *CNN* on June 23, 2015, entitled “‘Emotional’ robot sells out in a ‘minute.’”
- <https://www.indiegogo.com/projects/jibo-the-world-s-first-social-robot-for-the-home#> — Cynthia Breazeal’s Indiegogo project page for “JIBO, the world’s first social robot for the home.”
- <https://www.economist.com/science-and-technology/2015/04/18/robochef-gets-cooking> — Article first appearing in *The Economist* on April 18, 2015, entitled “Robochef gets cooking: The ultimate kitchen gadget for the home that has everything: A robotic chef.”
- http://www.humanengineers.com/hr_library/the-changing-face-of-business-and-the-part-artificial-intelligence-has-to-play — Article from the World Economic Forum, first appearing on January 11, 2017, entitled “The changing face of business—and the part artificial intelligence has to play.”
- <https://www.youtube.com/watch?v=D4NPQ8mfKU0> — Video of 1986 Pixar short film, *Luxo, Jr.*, written and directed by John Lasseter.
- <https://www.youtube.com/watch?v=qy02lwwGv3U> — Video of Hoffman and his robotic marimba’s improvisational take on Duke Jordan’s *Jordu*; April 22, 2009.
- http://www.chinadaily.com.cn/business/tech/2017-01/27/content_28065042.htm — Article by Fan Feifei, first appearing in *China Daily* on January 27, 2017, entitled “Will artificial companions be our best friend in the future?”