Bibliography

Adams, E. W. (1975). *The Logic of Conditionals: An Application of Probability to Deductive Logic*, volume 86. Springer Science & Business Media.

Ajtai, M. (1996). Generating hard instances of lattice problems. In *Proceedings of the 28th annual ACM symposium on Theory of computing*, pages 99–108.

Ajtai, M. and Dwork, C. (1997). A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 284–293.

Alekhnovich, M., Hirsch, E. A., and Itsykson, D. (2005). Exponential lower bounds for the running time of DPLL algorithms on satisfiable formulas. *Journal of Automated Reasoning*, 35(1):51–72.

Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843.

Allen, S. R., O'Donnell, R., and Witmer, D. (2015). How to refute a random CSP. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 689–708. IEEE.

Ansótegui, C., Bonet, M. L., Levy, J., and Manyá, F. (2008). Measuring the hardness of SAT instances. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 222–228.

Anthony, M. and Bartlett, P. (1999). *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge.

Applebaum, B., Barak, B., and Xiao, D. (2008). On basing lower-bounds for learning on worst-case assumptions. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 211–220. IEEE.

Arden, B. W., Galler, B. A., and Graham, R. M. (1961). An algorithm for equivalence declarations. *Communications of the ACM*, 4(7):310–314.

Atserias, A. and Bonet, M. L. (2004). On the automatizability of resolution and related propositional proof systems. *Information and Computation*, 189(2):182–201.

Atserias, A. and Müller, M. (2020). Automating resolution is NP-hard. *Journal of the ACM*, 67(5):1–17.

Audibert, J.-Y. and Tsybakov, A. B. (2007). Fast learning rates for plug-in classifiers. *The Annals of Statistics*, pages 608–633.

Awasthi, P., Blum, A., and Sheffet, O. (2010). Improved guarantees for agnostic learning of disjunctions. In *Proceedings of the 23rd Conference on Learning Theory, COLT 2010*, pages 359–367. Omnipress.

Baader, F. and Snyder, W. (2001). Unification theory. In Robinson, A. J. A. and Voronkov, A., editors, *Handbook of Automated Reasoning*, pages 445–533. Elsevier.

Bacchus, F., Grove, A. J., Halpern, J. Y., and Koller, D. (1996). From statistical knowledge bases to degrees of belief. *Artificial intelligence*, 87(1-2):75–143.

Banaszczyk, W. (1993). New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635.

Bayardo Jr., R. J. and Schrag, R. C. (1997). Using CSP look-back techniques to solve real-world SAT instances. In *Proceedings of the 14th Nationall Conference on Artificial Intelligence (AAAI'97)*, pages 203–208.

Beame, P., Kautz, H., and Sabharwal, A. (2004). Towards understanding and harnessing the potential of clause learning. *Journal of Artificial Intelligence Research*, 22:319–351.

Beame, P. and Pitassi, T. (1996). Simplified and improved resolution lower bounds. In *Proceedings of the 37th Conference on Foundations of Computer Science*, pages 274–282.

Belle, V. (2017). Open-universe weighted model counting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Belle, V. and Juba, B. (2019). Implicitly learning to reason in first-order logic. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, pages 3381–3391. Curran Associates, Inc.

Belson, W. A. (1959). Matching and prediction on the principle of biological classification. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 8(2):65–75.

Ben-Sasson, E. and Wigderson, A. (2001). Short proofs are narrow – resolution made simple. *Journal of the ACM*, 48(2):149–169.

Bernoulli, J. (1713). *Are Conjectandi*. Impensis Thurnisiorum, fratrum.

Biere, A., van Maaren, H., and Walsh, T. (2021). *Handbook of Satisfiability: Second Edition*. Frontiers in Artificial Intelligence and Applications (IOS Press). SAGE Publications.

Blake, A. (1937). *Canonical expressions in Boolean algebra*. PhD thesis, The University of Chicago.

Blum, A. (1992). Rank-r decision trees are a subclass of r-decision lists. *Information Processing Letters*, 42:183–185.

Blum, A., Furst, M., Kearns, M., and Lipton, R. J. (1994). Cryptographic primitives based on hard learning problems. In Stinson, D. R., editor, *Proceedings of the 13th Annual International Cryptology Conference on Advances in Cryptology – CRYPTO'93*, volume 773 of *Lecture Notes in Computer Science*, pages 278–291. Springer-Verlag.

Blum, M. (1982). Coin flipping by telephone. In *Proceedings of the 24th IEEE Computer Conference (COMPCON)*, pages 133–137. See also SIGACT News volume 15, number 1, 1983.

Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M. K. (1987). Occam's razor. *Information Processing Letters*, 24:377–380.

Blumer, A., Ehrenfeucht, A., Haussler, D., and Warmuth, M. K. (1989). Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):929–965.

Bonet, M. L., Domingo, C., Gavaldá, R., Maciel, A., and Pitassi, T. (2004). Non-automatizability of bounded-depth Frege proofs. *Computational Complexity*, 13:47–68.

Bonet, M. L., Pitassi, T., and Raz, R. (2000). On interpolation and automization for Frege proof systems. *SIAM Journal on Computing*, 29(6):1939–1967.

Boole, G. (1847). *The Mathematical Analysis of Logic*. Philosophical Library, Incorporated.

Boole, G. (1854). *An investigation of the laws of thought: on which are founded the mathematical theories of logic and probabilities*. Walton and Maberly.

Bousquet, O., Hanneke, S., Moran, S., Shafer, J., and Tolstikhin, I. (2023). Fine-grained distribution-dependent learning curves. In Neu, G. and Rosasco, L., editors, *Proceedings of*

Thirty Sixth Conference on Learning Theory, volume 195 of *Proceedings of Machine Learning Research*, pages 5890–5924. PMLR.

Bousquet, O., Hanneke, S., Moran, S., Van Handel, R., and Yehudayoff, A. (2021). A theory of universal learning. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 532–541.

Brachman, R. and Levesque, H. (2004). *Knowledge representation and reasoning*. Elsevier.

Bshouty, N. H. and Burroughs, L. (2005). Maximizing agreements with one-sided error with applications to heuristic learning. *Machine Learning*, 59(1–2):99–123.

Bylander, T., Allemang, D., Tanner, M. C., and Josephson, J. R. (1991). The computational complexity of abduction. *Artificial Intelligence*, 49:25–60.

Carlson, A., Betteridge, J., Kisiel, B., Settles, B., Hruschka, E., and Mitchell, T. (2010). Toward an architecture for never-ending language learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 24, pages 1306–1313.

Carnap, R. (1950). *Logical Foundations of Probability*. University of Chicago Press, Chicago, IL.

Carroll, L. (1897). *Symbolic Logic: Part 1, Elementary*. Macmillan.

Charniak, E. and McDermott, D. (1985). *Introduction to Artificial Intelligence*. Addison-Wesley, Reading, MA.

Church, A. (1936). A note on the entscheidungsproblem. *The Journal of Symbolic Logic*, 1(1):40–41.

Clark, A. (1987). From folk psychology to naive psychology. *Cognitive Science*, 11(2):139–154.

Cobham, A. (1965). The intrinsic computational difficulty of functions. In Bar-Hillel, Y., editor, *Logic, methodology and philosophy of science*, pages 24–30. North-Holland Pub. Co.

Codd, E. F. (1970). A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387.

Cohn, A. G. (1995). The challenge of qualitative spatial reasoning. *ACM Computing Surveys (CSUR)*, 27(3):323–325.

Cohn, D. and Tesauro, G. (1992). How tight are the Vapnik-Chervonenkis bounds? *Neural Computation*, 4(2):249–269.

Cook, S. A. (1971). The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158.

Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2022). *Introduction to Algorithms*. MIT Press, fourth edition.

Cox, P. T. and Pietrzykowski, T. (1986). Causes for events: their computation and applications. In *Proceedings of the 8th Internationall Conference on Automated Deduction*, pages 608–621.

Craig, W. (1957). Three uses of the herbrand-gentzen theorem in relating model theory and proof theory. *The Journal of Symbolic Logic*, 22(3):269–285.

Dalal, M. and Etherington, D. W. (1992). A hierarchy of tractable satisfiability problems. *Information Processing Letters*, 44(4):173–180.

Daniely, A., Linial, N., and Shalev-Shwartz, S. (2014). From average case complexity to improper learning complexity. In *Proceedings of the 46th Symposium on Theory of Computing*, pages 441–448.

Daniely, A. and Shalev-Shwartz, S. (2016). Complexity theoretic limitations on learning DNF's. In *Proceedings of the 29th Conference on Computational Learning Theory*,

volume 49 of *JMLR Workshops & Conference Proceedings*, pages 1–16.

Davis, M., Logemann, G., and Loveland, D. W. (1962). A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397.

Davis, M. and Putnam, H. (1960). A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215.

De Morgan, A. (1847). *Formal Logic: Or, the Calculus of Inference, Necessary and Probable*. Taylor and Waldon.

De Rezende, S. F., Göös, M., Nordström, J., Pitassi, T., Robere, R., and Sokolov, D. (2021). Automating algebraic proof systems is NP-hard. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 209–222.

Diffie, W. and Hellman, M. (1976). New directions in cryptography. *IEEE Transactions on Information Theory*, 22:423–439.

Dirichlet, G. L. (1842). Verallgemeinerung eines Satzes aus der Lehre von den Kettenbrüchen nebst einigen Anwendungen auf die Theorie der Zahlen. *SB Preuss. Akad. Wiss.*, pages 93–95. Reprinted in L. Kronecker (ed.), G. L. Dirichlet's Werke volume I, pages 635–638, Chelsea, New York, 1969.

Dowling, W. F. and Gallier, J. H. (1984). Linear-time algorithms for testing the satisfiability of propositional horn formulae. *The Journal of Logic Programming*, 1(3):267–284.

Durgin, A. and Juba, B. (2019). Hardness of improper one-sided learning of conjunctions for all uniformly falsifiable CSPs. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory*, volume 98 of *Proceedings of Machine Learning Research*, pages 369–382.

Edmonds, J. (1965). Paths, trees, and flowers. *Canadian Journal of mathematics*, 17:449–467.

Ehrenfeucht, A. and Haussler, D. (1989). Learning decision trees from random examples. *Information and Computation*, 82(3):231–246.

Ehrenfeucht, A., Haussler, D., Kearns, M., and Valiant, L. (1989). A general lower bound on the number of examples needed for learning. *Information and Computation*, 82(3):247–261.

Esteban, J. L. and Torán, J. (2001). Space bounds for resolution. *Information and Computation*, 171(1):84–97.

Feldman, V., Gopalan, P., Khot, S., and Ponnuswami, A. K. (2009). On agnostic learning of parities, monomials, and halfspaces. *SIAM Journal on Computing*, 39(2):606–645. Preliminary versions appeared in 47th FOCS, 2006 and 21st CCC, 2006.

Fikes, R. E. and Nilsson, N. J. (1971). STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3–4):189–208.

Filmus, Y., Lauria, M., Nordstrom, J., Ron-Zewi, N., and Thapen, N. (2015). Space complexity in polynomial calculus. *SIAM Journal on Computing*, 44(4):1119–1153.

Finger, J. J. (1987). *Exploiting constraints in design synthesis*. PhD thesis, Stanford University.

Frege, G. (1879). *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle.

Freund, Y. (1995). Boosting a weak learning algorithm by majority. *Information and Computation*, 121(2):256–285.

Freund, Y. and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139.

Galil, Z. (1977). On resolution with clauses of bounded size. *SIAM Journal on Computing*, 6:444–459.

Galler, B. A. and Fisher, M. J. (1964). An improved equivalence algorithm. *Communications of the ACM*, 7(5):301–303.

Gentzen, G. (1936). Die widerspruchsfreiheit der reinen zahlentheorie. *Mathematische annalen*, 112(1):493–565.

Gödel, K. (1930). Die vollständigkeit der axiome des logischen funktionenkalküls. *Monatshefte für Mathematik und Physik*, 37(1):349–360.

Goldman, S. A., Kwek, S. S., and Scott, S. D. (2003). Learning from examples with unspecified attribute values. *Information and Computation*, 180(2):82–100.

Goldreich, O. (2001). *Foundations of cryptography: volume 2, basic applications*. Cambridge University Press.

Goldreich, O., Goldwasser, S., and Micali, S. (1986). How to construct random functions. *Journal of the ACM*, 33(4):792–807.

Gomes, C. P., Selman, B., and Crato, N. (1997). Heavy-tailed distributions in combinatorial search. In *Proceedings of the 3rd International Conference on Principles and Practice of Constraint Programming (CP97)*, volume 1330 of *LNCS*, pages 121–135. Springer.

Green, C. (1969). Application of theorem proving to problem solving. In *Proceedings of the 1st International Joint Conference on Artificial intelligence*, pages 219–239.

Grosof, B. N. (1986). An inequality paradigm for probabilistic knowledge: the logic of conditional probability intervals. In *Machine Intelligence and Pattern Recognition*, volume 4, pages 259–275. Elsevier.

Hacking, I. (1975). *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction and Statistical Inference*. Cambridge University Press.

Haken, A. (1985). The intractability of resolution. *Theoretical computer science*, 39:297–308.

Hanks, S. and McDermott, D. (1987). Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33(3):379–412.

Hanneke, S. (2016). The optimal sample complexity of PAC learning. *Journal of Machine Learning Research*, 17(38):1–15.

Hanneke, S. and Moran, S. (2026). A theory of universal agnostic learning. arXiv preprint arXiv:2601.20961.

Harchol-Balter, M. (2023). *Introduction to probability for computing*. Cambridge University Press.

Hartmanis, J. and Stearns, R. E. (1965). On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306.

Håstad, J. (1999). Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, 182(1):105–142.

Håstad, J., Impagliazzo, R., Levin, L. A., and Luby, M. (1999). A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396.

Haussler, D. (1989). Learning conjunctive concepts in structural domains. *Machine Learning*, 4(1):7–40.

Haussler, D., Kearns, M., Littlestone, N., and Warmuth, M. K. (1991). Equivalence of models for polynomial learnability. *Information and Computation*, 95(2):129–161.

Haviv, I. and Regev, O. (2007). Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 469–477.

Hayes, P. J. (1979). The naive physics manifesto. In *Expert systems in the microelectronic age*. Edinburgh University Press.

Hirahara, S. (2022). NP-hardness of learning programs and partial MCSP. In *2022 IEEE 63rd annual symposium on foundations of computer science (FOCS)*, pages 968–979. IEEE.

Hobbs, J. R., Stickel, M. E., Appelt, D. E., and Martin, P. (1990). Interpretation as abduction. Technical Report 499, SRI, Menlo Park, CA.

Hopcroft, J. E. and Ullman, J. D. (1971). A linear list merging algorithm. Technical Report CU-CSD-71-111, Cornell University.

Horton, R. E. (1945). Erosioned development of streams and their drainage basins, hydrophysical approach to quantitative morphology. *Bull. Geol. Soc. of America*, 56:275–370.

Huet, G. (1976). *Résolution d'Équations dans des Langages d'Order 1, 2, . . . ω*. PhD thesis, Universite de Paris VII.

Hume, D. (1748). *Philosophical Essays Concerning Human Understanding*. A. Millar. Later retitled "An Enquiry Concerning Human Understanding".

Jevons, W. S. (1874). *The Principles of Science: A Treatise on Logic and Scientific Method*, volume 1 of *The Principles of Science: A Treatise on Logic and Scientific Method*. Macmillan and Company.

Juba, B. (2012). Learning implicitly in reasoning in PAC-semantics. arXiv preprint arXiv:1209.0056.

Juba, B. (2013). Implicit learning of common sense for reasoning. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 939–946.

Juba, B. (2016a). Integrated common sense learning and planning in POMDPs. *Journal of Machine Learning Research*, 17(96):1–37.

Juba, B. (2016b). Learning abductive reasoning using random examples. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 999–1007.

Juba, B. and Le, H. S. (2019). Precision-recall versus accuracy and the role of large data sets. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pages 4039–4048.

Juba, B. and Li, H. (2020). More accurate learning of k-DNF reference classes. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, volume 4, pages 4385–4393.

Juba, B., Li, Z., and Miller, E. (2018). Learning abduction under partial observability. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 1888–1896.

Jukna, S. (2012). *Boolean function complexity: advances and frontiers*, volume 27. Springer.

Kalai, A. T., Kanade, V., and Mansour, Y. (2012). Reliable agnostic learning. *Journal of Computing and System Sciences*, 78:1481–1495.

Kanade, V. and Thaler, J. (2014). Distribution-independent reliable learning. In *Proceedings of the 27th Conference on Computational Learning Theory*, volume 35 of *JMLR Workshops and Conference Proceedings*, pages 3–24.

Kautz, H., Kearns, M., and Selman, B. (1995). Horn approximations of empirical data. *Artificial intelligence*, 74(1):129–145.

Kautz, H. and Selman, B. (1992). Planning as satisfiability. In *Proceedings of the 10th European Conference on Artificial Intelligence (ECAI-92)*, pages 359–363.

Kautz, H. and Selman, B. (1996). Pushing the envelope: Planning, propositional logic, and stochastic search. In *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96)*, pages 1194–1201.

Kearns, M. and Valiant, L. (1994). Cryptographic limitations on learning Boolean formulae and finite automata. *Journal of the ACM*, 41:67–95.

Kearns, M. J., Schapire, R. E., and Sellie, L. M. (1994). Towards efficient agnostic learning. *Machine Learning*, 17(2-3):115–141.

Kearns, M. J. and Vazirani, U. (1994). *An introduction to computational learning theory*. MIT press.

Khardon, R. and Roth, D. (1997). Learning to reason. *Journal of the ACM*, 44(5):697–725.

Khot, S. (2005). Hardness of approximating the shortest vector problem in lattices. *Journal of the ACM*, 52(5):789–808.

Kleinberg, J. and Tardos, É. (2006). *Algorithm Design*. Pearson/Addison-Wesley.

Klivans, A. R. and Sherstov, A. A. (2009). Cryptographic hardness for learning intersections of halfspaces. *Journal of Computer and System Sciences*, 75(1):2–12.

Koch, C., Strassle, C., and Tan, L.-Y. (2023). Superpolynomial lower bounds for decision tree learning and testing. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1962–1994. SIAM.

Koltchinskii, V. and Beznosova, O. (2005). Exponential convergence rates in classification. In Auer, P. and Meir, R., editors, *Learning Theory (COLT 2005)*, number 3559 in Lecture Notes in Artificial Intelligence, pages 295–307. Springer, Berlin, Heidelberg.

Kowalski, R. (1974). Predicate logic as programming language. In *Proceedings of the International Federation of Information Processing Congress*, pages 569–544. North-Holland.

Kozen, D. C. (2012). *The Design and Analysis of Algorithms*. Monographs in Computer Science. Springer.

Krajíček, J. (1994). Lower bounds to the size of constant-depth propositional proofs. *The Journal of Symbolic Logic*, 59(1):73–86.

Krajíček, J. (2019). *Proof complexity*, volume 170. Cambridge University Press.

Krajíček, J. and Pudlák, P. (1995). Some consequences of cryptographical conjectures for S_2^1 and EF. In Leivant, D., editor, *Logic and Computational Complexity*, volume 960 of *LNCS*, pages 210–220. Springer, Berlin.

Kullmann, O. (1999). Investigating a general hierarchy of polynomially decidable classes of CNF's based on short tree-like resolution proofs. Technical Report TR99-041, ECCC.

Kyburg, H. E. (1974). *The Logical Foundations of Statistical Inference*. Reidel, Dordrecht.

Lagrange, J. L. (1773). Recherches d'arithmétique. *Nouveaux Mémoires de l'Académie royale des Sciences et Belles-Lettres de Berlin*, pages 265–312.

Lakemeyer, G. and Levesque, H. J. (2002). Evaluation-based reasoning with disjunctive information in first-order knowledge bases. In *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning*, pages 73–81.

Lenat, D. B. (1995). CYC: a large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38.

Lenstra, A. K., Lenstra Jr, H. W., and Lovász, L. (1982). Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534.

Levesque, H. J. (1998). A completeness result for reasoning with incomplete first-order knowledge bases. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning*, pages 14–23.

Levin, L. A. (1973). Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266.

Li, J. and Nguyen, P. Q. (2025). A complete analysis of the BKZ lattice reduction algorithm. *Journal of Cryptology*, 38(1):12.

Liu, Y., Lakemeyer, G., and Levesque, H. J. (2004). A logic of limited belief for reasoning with disjunctive information. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning*, pages 587–597.

Livni, R., Shalev-Shwartz, S., and Shamir, O. (2014). On the computational efficiency of training neural networks. In *Advances in neural information processing systems 27*, pages

855–863.

Lyubashevsky, V. and Micciancio, D. (2009). On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In *Annual International Cryptology Conference*, pages 577–594. Springer.

Maciel, A. and Pitassi, T. (1998). Towards lower bounds for bounded-depth Frege proofs with modular connectives. In Beame, P. and Buss, S., editors, *Proof Complexity and Feasible Arithmetics*, number 39 in DIMACS Ser. Discrete Math. Theoret. Comput. Sci., pages 195–227. AMS.

Marques-Silva, J. P. and Sakallah, K. A. (1999). GRASP: a search algorithm for propositional satisfiability. *IEEE Transactions on Computing*, 48(5):506–521.

Martelli, A. and Montanari, U. (1982). An efficient unification algorithm. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(2):258–282.

McCarthy, J. (1959). Programs with common sense. In *Teddington Conf. on the Mechanization of Thought Processes*, pages 756–791. Available at http://www-formal.stanford.edu/jmc/mcc59.html.

McCarthy, J. (1977). Epistemological problems of artificial intelligence. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence*, pages 1038–1044. Morgan Kaufmann.

McCarthy, J. (1980). Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1–2):27–39. Available at http://www-formal.stanford.edu/jmc/circumscription.html.

McCarthy, J. (1988). Mathematical logic in artificial intelligence. *Daedalus*, 117(1):297–311.

McCarthy, J. (1998). Elaboration tolerance. In *Common sense*.

McCarthy, J. and Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, B. and Michie, D., editors, *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press.

McDermott, D. and Doyle, J. (1980). Non-monotonic logic I. *Artificial intelligence*, 13(1-2):41–72.

Micciancio, D. and Peikert, C. (2012). Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proceedings of the 31st Annual International Conference on Theory and Applications of Cryptographic Techniques*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718. Springer Berlin Heidelberg.

Michael, L. (2009). Reading between the lines. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 1525–1530.

Michael, L. (2010). Partial observability and learnability. *Artificial Intelligence*, 174(11):639–669.

Michael, L. (2013). Machines with websense. In *Proceedings of 11th International Symposium on Logical Formalizations of Commonsense Reasoning (Commonsense'13)*.

Michael, L. (2014). Simultaneous learning and prediction. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning*, pages 348–357.

Mill, J. S. (1843). *A System of Logic, Ratiocinative and Inductive*, volume 1. John W. Parker, London.

Minkowski, H. (1896). *Geometrie der Zahlen*. Teubner.

Minsky, M. (1975). A framework for representing knowledge. In Winston, P., editor, *The psychology of computer vision*, pages 211–277. McGraw-Hill.

Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2018). *Foundations of machine learning*. MIT Press.

Mueller, E. T. (2014). *Commonsense Reasoning: An Event Calculus Based Approach*. Morgan Kaufmann.

Nakashima, A. (1935). The theory of relay circuit composition. *Journal of the Institute of Telegraph and Telephone Engineers of Japan*, 150:731–752.

Nitanda, A. and Suzuki, T. (2019). Stochastic gradient descent with exponential convergence rates of expected classification errors. In Chaudhuri, K. and Sugiyama, M., editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1417–1426. PMLR.

Papadimitriou, C. H. (1994). *Computational Complexity*. Addison-Wesley.

Paterson, M. S. and Wegman, M. N. (1978). Linear unification. *Journal of Computer and System Sciences*, 16(2):158–167.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

Peikert, C. (2009). Public-key cryptosystems from the worst-case shortest vector problem. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 333–342.

Peirce, C. S. (1865). On an improvement in Boole's calculus of logic. In *Proceedings of the American Academy of Arts and Sciences*, volume 7, pages 249–261. American Academy of Arts Sciences.

Peirce, C. S. (1870). *Description of a notation for the logic of relatives, resulting from an amplification of the conceptions of Boole's calculus of logic*. Welch, Bigelow.

Peirce, C. S. (1885). On the algebra of logic: A contribution to the philosophy of notation. *American Journal of Mathematics*, 7(2-3):180–196.

Peirce, C. S. (1931). Elements of logic. In Hartshorn, C., et al., editor, *Collected Papers of Charles Sanders Peirce*. Harvard University Press.

Pillaud-Vivien, L., Rudi, A., and Bach, F. (2018). Exponential convergence of testing error for stochastic gradient methods. In Bubeck, S., Perchet, V., and Rigollet, P., editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 250–296. PMLR.

Pitt, L. and Valiant, L. G. (1988). Computational limitations on learning from examples. *Journal of the ACM*, 35(4):965–984.

Pitt, L. and Warmuth, M. K. (1990). Prediction-preserving reducibility. *Journal of Computer and System Sciences*, 41(3):430–467.

Pollock, J. L. (1990). *Nomic Probabilities and the Foundations of Induction*. Oxford University Press, Oxford.

Poole, D. (1990). A methodology for using a default and abductive reasoning system. *International Journal of Intelligent Systems*, 5:521–548.

Poole, D. (1993). Probabilistic Horn abduction and Bayesian networks. *Artificial Intelligence*, 64(1):81–129.

Pople, H. (1973). On the mechanization of abductive logic. In *Proceedings of the 3rd International Joint Conference on Artificial Intelligence*, pages 147–152.

Post, E. L. (1921). Introduction to a general theory of elementary propositions. *American journal of mathematics*, 43(3):163–185.

Pudlák, P. (2000). Proofs as games. *The American Mathematical Monthly*, 107(6):541–550.

Pudlák, P. and Buss, S. R. (1995). How to lie without being (easily) convicted and the lengths of proofs in propositional calculus. In *Computer Science Logic: 8th Workshop, CSL'94 Kazimierz, Poland, September 25–30, 1994 Selected Papers 8*, pages 151–162. Springer.

Qi, D., Arfin, J., Zhang, M., Mathew, T., Pless, R., and Juba, B. (2018). Anomaly explanation using meta-data. In *Proceedings of the 2018 IEEE Winter Conference on Applications of Computer Vision*, pages 1916–1924.

Quine, W. V. O. (1960). *Word and Object*. MIT Press, Cambridge, MA.

Quine, W. V. O. (1969). Natural kinds. In *Ontological Relativity and Other Essays*. Columbia University Press.

Reckhow, R. A. (1976). *On the lengths of proofs in the propositional calculus*. PhD thesis, University of Toronto.

Regev, O. (2009). On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):1–40.

Reichenbach, H. (1949). *Theory of Probability*. University of California Press, Berkeley, CA.

Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13:81–132.

Reiter, R. and Criscuolo, G. (1983). Some representational issues in default reasoning. In *Computational Linguistics*, pages 15–27. Elsevier.

Reiter, R. and de Kleer, J. (1987). Foundations for assumption-based truth maintenance systems: Preliminary report. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, pages 183–188.

Ribeiro, M. T., Singh, S., and Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*, pages 1527–1535.

Rivest, R. L. (1987). Learning decision lists. *Machine Learning*, 2(3):229–246.

Robinson, J. A. (1965). A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41.

Roth, D. (1995). Learning to reason: the non-monotonic case. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, volume 2, pages 1178–1184.

Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3):581–592.

Samson, E. W. and Mueller, R. K. (1955). Circuit minimization: sum to one process for irredundant sums. Technical Report 55-118, Communications Laboratory, Electronics Research Directorate, Air Force Cambridge Research Center, Cambridge, MA.

Sauer, N. (1972). On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147.

Schapire, R. E. (1990). The strength of weak learnability. *Machine Learning*, 5(2):197–227.

Schapire, R. E. and Freund, Y. (2012). *Boosting: Foundations and Algorithms*. MIT Press.

Schnorr, C.-P. (1987). A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical computer science*, 53(2-3):201–224.

Schnorr, C.-P. and Euchner, M. (1994). Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 66(1):181–199.

Schuurmans, D. (1997). Characterizing rational versus exponential learning curves. *Journal of Computer and System Sciences*, 55(1):140–160.

Shalev-Shwartz, S. and Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, New York, NY.

Shalev-Shwartz, S., Shamir, O., and Sridharan, K. (2011). Learning kernel-based halfspaces with the 0-1 loss. *SIAM Journal on Computing*, 40(6):1623–1646.

Shannon, C. E. (1938). A symbolic analysis of relay and switching circuits. *Transactions of the American Institute of Electrical Engineers*, 57(12):713–723.

Shelah, S. (1972). A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41(1):247–261.

Shestakov, V. I. (1938). *Some mathematical methods for construction and simplification of two terminal electrical networks of class A*. PhD thesis, Lomonosov State University.

Shor, P. W. (1994). Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134. IEEE Computer Society.

Sipser, M. (2012). *Introduction to the Theory of Computation*. Cengage Learning.

Skolem, T. (1928). Über die mathematische logik. *Norsk matematisk tidsskrift*, 10:125–142.

Spelke, E. S. and Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10(1):89–96.

Spira, P. (1971). On time-hardware complexity tradeoffs for Boolean functions. In *Proceedings of the 4th Hawaii Symposium on System Sciences*, pages 525–527.

Stallman, R. M. and Sussman, G. J. (1977). Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis. *Artificial intelligence*, 9(2):135–196.

Stork, D. G. (1999). The open mind initiative. *IEEE Expert Systems and Their Applications*, 14(3):16–20.

Strahler, A. N. (1952). Hypsometric (area-altitude) analysis of erosional topology. *Bull. Geol. Soc. of America*, 63:1117–1142.

Talagrand, M. (1994). Sharper bounds for gaussian and empirical processes. *The Annals of Probability*, 22:28–76.

Tarjan, R. E. (1975). Efficiency of a good but not linear set union algorithm. *Journal of the ACM*, 22(2):215–225.

Tarski, A. (1956). The concept of truth in formalized languages. In *Logic, semantics, metamathematics*, pages 152–278. Clarendon Press.

Tseitin, G. S. (1968). On the complexity of derivation in propositional calculus. In *Structures in Constructive Mathematics and Mathematical Logic*, pages 115–125. Steklov Mathematical Institute.

Turing, A. M. (1936). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(3-4):230–265.

Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, 18(11):1134–1142.

Valiant, L. G. (1985). Learning disjunction of conjunctions. In *Proceedings of the 9th International Joint Conference on Artificial intelligence*, volume 1, pages 560–566.

Valiant, L. G. (1994). *Circuits of the Mind*. Oxford University Press, Oxford.

Valiant, L. G. (1995). Rationality. In *Proceedings of the 8th Conference on Computational Learning Theory*, pages 3–14.

Valiant, L. G. (2000a). A neuroidal architecture for cognitive computation. *Journal of the ACM*, 47(5):854–882.

Valiant, L. G. (2000b). Robust logics. *Artificial Intelligence*, 117:231–253.

Valiant, L. G. (2006). Knowledge infusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1546–1551.

Vapnik, V. and Chervonenkis, A. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280.

Watt, S. (1995). A brief naive psychology manifesto. *Informatica*, 19(4):495–500.

Wittgenstein, L. (1922). *Tractatus Logico-philosophicus*. International library of psychology, philosophy, and scientific method. Harcourt, Brace, Incorporated.

Wittgenstein, L. (1953). *Philosophical Investigations*. Macmillan.

Zhang, L., Madigan, C. F., Moskewicz, M. W., and Malik, S. (2001). Efficient conflict driven learning in a Boolean satisfiability solver. In *Proceedings of the IEEE/ACM Internationall Conference on Computer Aided Design (ICCAD'01)*, pages 279–285.

Zhang, M., Mathew, T., and Juba, B. (2017). An improved algorithm for learning to perform exception-tolerant abduction. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 1257–1265.