# **Ernest Davis**

Contact:	Courant Institute of Mathematical Sciences	
	251 Mercer St.	
	New York, NY 10012	
	Phone: (212) 998-3123	
	Fax: (212) 995-4121	
	Email: DAVISE@CS.NYU.EDU	
	Web: http://cs.nyu.edu/faculty/davise/index.html	
Education:	Yale University Ph.D. (Computer Science) 1984	
	MIT B.Sc. (Mathematics) 1977	
Employment:	Computer Science Dept., New York University	
	Professor	2008 -
	Associate Professor	1989 - 2008
	Assistant Professor	1983–1989
	Computer Science Dept., Yale University	
	Research Assistant	1980 - 83
	Teaching Assistant	1980-82
	Computer Graphics Dept., CE Lummus Inc., Bloomfield, NJ	
	Program Analyst	1977 - 79
	Math Dept. M.I.T.	
	Calculus Tutor	1975 - 77.
Honors:	Keynote Speaker, COSIT-11	2011
	Outstanding Program Committee Member, AAAI-06	2006
	Knowledge Representation and Reasoning Distinguished Lecturer, Universities of York and Leeds.	1996
	IBM Graduate Fellowship	1982 - 83
	NSF Graduate Fellowship	1979 - 82
	Phi Beta Kappa	1977
	Sigma Xi	1977
	Putnam Fellow, Putnam Mathematical Test	1976

# **Professional Activities**

# Journal editorships

Associate editor, Artificial Intelligence Journal, 2016-2020. Area editor, ACM Transactions on Computational Logic, 2008-2014. Guest Editor (with L. Morgenstern), Artificial Intelligence vol. 153, nos. 1-2, special issue on Logical Formalizations of Commonsense Reasoning, March 2004. Book reviews editor. IEEE Ernert, 1004 1008

Book reviews editor, IEEE Expert, 1994-1998.

#### Journal reviewing

AI Journal, JACM, International Journal of Approximate Reasoning, IEEE Transactions on Software Engineering, Cognitive Science, IEEE Transactions on Knowledge and Data Engineering, Annals of Mathematics and Artificial Intelligence, IEEE Transactions on Pattern Analysis and Machine Intelligence, Journal of Logic and Computation, ACM Computing Surveys, Computational Intelligence, Fundamenta Informaticae, Mathematical Reviews, Spatial Cognition and Computation, Computing Reviews, Journal of Philosophical Logic, Program, Comm. ACM, Applied Ontology. Behavioral and Brain Sciences, Science, Proc. National Academy of Sciences (PNAS), Information Systems, Minds and Machines, IEEE Transactions on Cognitive and Developmental Systems, International Journal of Geographical Information Science, Journal of Experimental Psychology: General, AI & Society, Trends in Cognitive Science, Knowledge-Based Systems, ACL Rolling Review, British Journal for the Philosophy of Science, World Wide Web, Digital Discovery, Artificial Intelligence Review, Journal of Intelligence.

#### Conference and workshop chair

Co-chair and organizer (with L. Morgenstern and K. Sanders), Workshop on Knowledge, Perception, and Planning, IJCAI-89.

Co-chair (with L. Morgenstern, J. McCarthy, and R. Reiter) of the Fifth Symposium on Logical Formalizations of Commonsense Reasoning, May 2001, New York City.

Co-chair (with P. Doherty and E. Erdem) of the Tenth Symposium on Logical Formalizations of Commonsense Reasoning, March 2011, Stanford University.

Co-chair (with Mehul Bhatt and Hans Guesgen) STeDy 2012 (International Workshop on Spatio-Temporal Dynamics).

Local Arrangements Chair: IJCAI-16.

Conference program committee or referee

Area Chair: KR-16, KR-18.

Senior Program Committee: AAAI-10, and committee on NECTAR papers at AAAI-10. IJCAI-11, AAAI-23.

Program Committee: IJCAI-87, Second Symposium on Logical Formalizations of Commonsense Reasoning 1992, KR-94, Math and AI 1996, KR-96, AAAI-97, KR-98, AAAI-98, Commonsense 98, Formal Ontologies for Intelligent Systems 1998, IJCAI-99, AAAI-99, BISFAI-99, AAAI-00, FOIS-01, IJCAI-01, KR-02, Commonsense-03, KR-04, FOIS-04, Commonsense-05, AAAI-06, FOIS-06, KR-06, Commonsense-07, AAAI-07, KR-08, AAAI-08, Commonsense-09, FOIS-10, KR-10, KR-12, AAAI-13, Commonsense-13, IJCAI-13, Qualitative Reasoning Workshop 2013, Advances in Cognitive Systems 2014, AAAI-14, AAAI-15, Commonsense-15, Beyond Turing (AAAI Workshop) 2015, IJCAI-15, Global Conference on Artificial Intelligence (GCAI) 2015, IJCAI-17, AAAI-17, ProSocrates 2017, Cognitum 2017, Advances in Cognitive Systems 2017, Commonsense 2017, AAAI-18, ACS-18, AAAI-19, IJCAI-19, ACS-19, CogSci-20, KR-20, ACS-20, ACS-21, AAAI-22, EBeM-22 (Evaluation Beyond Metrics), ECAI-23.

Referee: IJCAI-85, IJCAI-89, IJCAI-91, IJCAI-93, IJCAI-95, IJCAI-05, Cognitive Science '05, IJCAI-07. STeDy 2010 (International Workshop on Spatio-Temporal Dynamics), STeDy 2012, CogSci-21, CogSci-22, CogSci-23, CogSci-24.

Rolling reviews for ACL (Association for Computational Linguistics).

Funding reviewing: NSF, NSERC, Allen Institute.

#### Other

Invited participant, Workshop on Mental Models, M.I.T., March 12-13, 1990.

Epistemic Logic and Its Applications (with L. Morgenstern), Tutorial, IJCAI-93.

Second reader for chapter 1 "First Order Logic" and chapter 9 "Qualitative Reasoning" in *The Handbook of Knowledge Representation*, Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter (eds.), Elsevier, Oxford, 2008.

Associate editor, ACM Computer Classification System, 2010-2011.

Reviewer, Grace Hopper Conference Scholarship proposals, 2015.

Organizing committee, Winograd Schema Challenge.

Beth Dissertation Prize Committee, 2017.

A Gentle Introduction to Deep Nets and Opportunities for the Future. Kenneth Church, Valia Kordoni, Gary Marcus, Ernest Davis, Yanjun Ma, Zeyu Chen. Tutorial, ACL 2022.

Participant, Cognitive Computational Neuroscience, General Adversarial Collaborations, "To what extent does the brain simulate the external world?", 2022.

Reviewer, Hans Sigrist Prize, 2023.

Membership: AAAI (Association for the Advancement of Artificial Intelligence).

# Ph.D. Students:

Leora Morgenstern, "Foundations of a Logic of Knowledge, Action, and Communication," September, 1988.

Leo Joskowicz, "Reasoning about Shape and Kinematic Function in Mechanical Devices," September, 1988.

Pasquale Caianiello, "Learning as the Evolution of Representation," November, 1989.

Alexander Botta, "A Theory of Natural Learning," May, 1991.

Jen-Lung Chiu, "Planning in an Imperfect World Using Previous Experiences," January 1995.

Tamir Klinger, "Adversarial Reasoning: A Logical Approach for Computer GO." January 2001.

Gedaminas Adomavicius, "Expert-Driven Validation of Set-Based Data Mining Results" (nominal co-advisor with Alex Tuzhilin.) July 2002.

Ji-Ae Shin, "TM-LPSAT: Encoding Temporal Metric Planning in Continuous Time," May 2004.

Tatiana Kichkaylo, "Construction of Component-Based Applications by Planning," (coadvisor with Vijay Karamcheti), December 2004

Ziyang Wang, "Incremental Web Search: Tracking Changes in the Web." May 2006.

Paul Bethe, "Advances in computer bridge: Techniques for a partial-information, communicationsbased game." January 2021.

#### MS theses supervised:

Yanai Lehavi, "Charlie: A Treaty Reinsurance Underwriter".

Kumar Shashi Prabh, "Performance of BLACKBOX Planning System on a Hard Problem of Satisfiability", May 2001.

Paul Bethe, "DTAC: A method for planning to claim in Bridge." May 2010.

Azam Asl, "A Qualitative Calculus for Three-Dimensional Rotations." (NYU Poly Computer Engineering Dept.) December 2011.

Emily Morton-Owens, "A tool for extracting and indexing spatio-temporal information from biographical articles in Wikipedia." May 2012.

Wei Peng, "Evaluating the commonsense reasoning abilities of pretrained language models," December 2019.

#### Senior theses supervised:

Yuling Gu, "Towards detecting temporal relations implicitly conveyed in text". May 2020.

Zhuoran (Jennifer) Zeng, "Physical reasoning in an open world". May 2022.

# Grants:

Reasoning about Shape and Function, NSF DCR-8402309, \$50,000, 12/15/84 - 5/31/86. Physical and Spatial Reasoning with Solid Objects, NSF DCR-8603758, \$82,600, 7/1/86 - 12/31/88.

Perception and Planning, (with L. Morgenstern), NSF IRI-8801529, \$115,000, 7/1/88 - 12/31/90.

Perception and Planning, (renewed, with no co-PI), NSF IRI-9001447, \$148,937, 8/1/90 - 7/31/93.

Knowledge Representation for Physical Reasoning, NSF IRI-9300446, \$180,000, 8/93-8/96.

Physical and Spatial Reasoning across Multiple Scales, NSF IRI-9625859, \$236,000, 8/96-8/99.

Commonsense Reasoning about Loosely Constrained Systems of Rigid Solid Objects, NSF IIS-0097537, \$289,000, 6/01-5/04.

Automating Commonsense Reasoning for Elementary Physical Science, NSF IIS-0534809, \$328,877, 2/06-8/10.

Adversarial Collaborative Research on Intuitive Physical Reasoning, NSF 2121102, co-PI with Todd Gureckis (Psychology, NYU) and Joshua Tenenbaum (Brain and Cognitive Science, MIT). \$329,00 total NYU; \$29,000 my part. 9/21-8/24.

#### **Invited Talks and Panel Participation:**

"Geographic Reasoning," Workshop on Naive Physics, University of Rochester, Spring 1982

"Planning and Execution in Navigation," DARPA Workshop on Planning and Robot Problem Solving, Washington D.C., May 15, 1986.

"A Logical Framework for Solid Object Physics," Workshop on Space Telerobotics, Jet Propulsion Labs, January 20-22, 1987; Workshop on Qualitative Physics, University of Illinois at Urbana, May 27-29, 1987; CIAR Graduate Student Workshop on Knowledge Representation, Edmonton, Alberta, June 6-8, 1988.

"Inferring Ignorance from the Locality of Visual Perception," IBM Watson Labs, July 1988.

"Error Correction in Cognitive Maps," SPIE Workshop on Sensor Fusion: Spatial Reasoning and Scene Interpretation, Boston, Nov. 19, 1988.

"Reasoning about Perception and Knowledge," University of Toronto, Jan. 19, 1989.

Chair, Panel on Temporal Reasoning, First International Conference on Principles of Knowledge Representation, May, 1989.

"Lucid Representations," *Bar Ilan Symposium on the Foundations of Artificial Intelligence*, Bar Ilan University, June 18, 1991 (Invited hour address); U. Connecticut at Storrs, November, 1991.

"QR work on Spatial Reasoning," IFIP Workshop on Knowledge Representation and Qualitative Reasoning, Islamorada, Fla, Feb. 1992.

Panel on "Hard Problems in Physical Reasoning," IFIP Workshop on Knowledge Representation and Qualitative Reasoning, Islamorada, Fla, Feb. 1992.

"Knowledge Preconditions for Plans," Bell Labs, April 22, 1993. IBM Watson Labs, June 1, 1993.

"Approximation and Abstraction in Solid Object Kinematics." Yale University, Dec. 1994. Rutgers University, Oct. 1995.

"The Automation of Commonsense Physical Reasoning" and "Qualitative Kinematics", Knowledge Representation and Reasoning Distinguished Lecturer, Universities of York and Leeds, May 1996.

"Formal Theories of Spatial Reasoning," NSF Workshop on Visual Cognition and Spatial Reasoning, Ellicott City, MD, May 15-17, 1997.

"Methodological Difficulties in Automating Commonsense Reasoning" Symposium on Architectures for Commonsense Reasoning, IBM Watson Labs, March 13-14, 2002.

"A First-Order Theory of Communication and Multi-Agent Plans," Computer Science Department Colloquium, City University of New York, November 11, 2004.

"Why Computers are So Stupid and What Can Be Done About It: Artificial Intelligence and Commonsense Knowledge," Palladium Lecture, Palladium Residence House, New York University, February 27, 2006.

"Commonsense Physical Reasoning: Boxes and Pitchers," IBM Watson Labs, June 6, 2007.

"Some metalogical properties of first-order languages that quantify over spatial regions." Seminar in Logic and Games, CUNY Graduate Center, October 8, 2010.

"Commonsense Reasoning about Chemistry Experiments: Ontologies and Representations", Commonsense-2009, Toronto, June 2, 2009; and University of Illinois at Chicago, April 12, 2011.

"Qualitative Spatial Reasoning in Interpreting Text and Narrative." Conference on Spatial Information Theory (Keynote address). September 13, 2011.

"Why Computers Are So Stupid and What Can Be Done About It," *Science on Saturdays*, Princeton Plasma Physics Lab, March 3, 2012.

Interviewed in *The Rise of Artificial Intelligence*, Off Book, PBS videos, produced by Lisa Romagnoli and Eric Brown, July 2013.

"The Scope and Limits of Simulation in Automated Reasoning and Cognitive Models," Seminar on Concepts and Categories, NYU Psychology Dept. March 28, 2014.

"How AI Programs Collect Concepts" Seminar on Concepts and Categories, NYU Psychology Dept. September 25, 2015.

"How Strong is the Empirical Evidence for Bayesian Models of Cognition?" with Gary Marcus. *Is the Brain Bayesian* symposium, NYU, December 4, 2015.

"Axiomatizing the Foundations of Physics, starting with the Experiments," Computationally Assisted Mathematical Discovery and Experimental Mathematics, May 14, 2016.

"Collecting Commonsense Inferences from Text." Cognitum 2016 (Workshop on Cognitive Knowledge Acquisition and Applications). July 11, 2016.

Moderator, Panel, "Progress and Caution in Artificial Intelligence", Blouin Creative Leadership Summit, September 2016.

"The Scope and Limits of Simulation in Automated Reasoning and Cognitive Models," Northwestern University, January 10, 2017. Cognitive AI Meetup, New York, February 23, 2017.

"Reasoning about Containers", Northwestern University, January 11, 2017. Renssalear Polytechnic, May 12, 2017.

Panel, "Getting Societal Benefits Right", SIROS-2 (Social Implications of Robotics Symposium), Brown University, March 30, 2017.

"The Logical Depth of Reasoning about Other Minds," Advances in Cognitive Systems, May 13, 2017.

"Proof Verification Technology and Physics," Google New York, Research Seminar, December 20, 2017,

"Four Challenges for Physical Reasoning," Army Research Lab, January 18, 2018.

"Building AIs with Common Sense," Princeton chapter of the ACM, May 16, 2019.

"Building Artificial Intelligence We Can Trust". Columbia University, class on Computers and Society. October 15, 2019. Columbia University, class on AI and Ethics, February 2020. MIT Task Force on the Work of the Future, MIT, May 2020. Renaissance Numérique, February 2021.

"Time and Space in Knowledge Graphs," Ontology Summit, May 2020.

"Using human skills taxonomies and tests as measures of artificial intelligence: Caveats" OECD, Expert Meeting on Skills and Tests for Assessing AI and Robotics, October 2020.

"The Scope and Limits of Simulation in Automated Physical Reasoning." Argonne National Labs, March 3, 2021.

"Automating Common Sense: Where do we stand?". Keynote address, Conference on Business Informatics, September 2, 2021.

"The Scope and Limits of Simulation in Automated Physical Reasoning." NIST, November 17, 2021.

"Common Sense and Artificial Intelligence," Cognitive Science Seminar, Rutgers, March 1, 2022.

"Artificial inteligence: Perceptions and Reality — Accomplishments, Challenges, Prospects, and Risks". Institute of International and European Affairs (Dublin). October 24, 2022

"Commonsense Physical Reasoning in Humans and Machines." Guest lecture, Mécanismes de l'intuition mathématique chez les êtres humains et les machines," Stanislas Dehaene, Collège de France, January 13, 2023

"Reconsidering *Rebooting*", University of Bamberg, May 15, 2023.

Panel, "Reliability of Current Large Language Models", NSF, September 5, 2023.

"AI and Elementary Science and Math Problems", Keynote talk, IEEE UEMCON conference, October 12, 2023.

"The Short Term Risks of Artificial Intelligence," Ethics Breakfast series, Dirah, January 21, 2024.

# **Publications**

Where not otherwise indicated, Ernest Davis is the sole author.

#### Books

- B.1 Representing and Acquiring Geographic Knowledge Pitman Press, London, 1986
- B.2 Representations of Commonsense Knowledge Morgan Kaufmann, San Mateo, CA, 1990.
- B.3 Linear Algebra and Probability for Computer Science Applications. CRC Press, A.K. Peters, 2012.
- B.4 Verses for the Information Age. Privately published. 2017.
- B.5 Rebooting AI: Building Artificial Intelligence We Can Trust. By G. Marcus and E. Davis. Pantheon Press. 2019. Translations into Chinese, Korean, and Russian.
- B.6 Particles, Fields, Measures, Stuff, and Things: Representations of Matter for Artificial Intelligence. In preparation.

# **Edited Books**

- E.1 Mathematics, Substance and Surmise: Views on the Meaning and Ontology of Mathematics,E. Davis and P. Davis (editors), Springer, 2015.
- E.2 "Will You Marry Me?" Some First-hand Accounts of Marriage Proposals, 1600-1900. Privately published, 2020.

#### Journal Articles

- J.1 Algorithms for Scheduling Tasks on Unrelated Processors. By E. Davis and J. Jaffe. JACM, Vol. 28 No. 4, October 1981, pp. 721-736
- J.2 What's the Point? By R. Schank, G. Collins, E. Davis, P. Johnson, S. Lytinen, and B. Reiser. Cognitive Science, Vol. 6, No. 3, 1982
- J.3 Planning and Executing Routes through Uncertain Territory. By D. McDermott and E. Davis. Artificial Intelligence, vol. 22, pp. 107-156, 1984
- J.4 Constraint Propagation with Interval Labels. Artificial Intelligence, vol. 32, 1987, pp. 281-331.

- J.5 A Logical Framework for Commonsense Predictions of Solid Object Behavior. AI in Engineering, vol. 3 no. 3, 1988, pp. 125-140.
- J.6 The Kinematics of Cutting Solid Objects. Annals of Mathematics and Artificial Intelligence, vol. 9, no. 3,4, 1993, pp. 253-305.
- J.7 Knowledge Preconditions for Plans. Journal of Logic and Computation, vol. 4, no. 5, Oct. 1994, pp. 721-766
- J.8 Order of Magnitude Comparisons of Distance Journal of AI Research, vol. 10, 1999, pp. 1-38.
- J.9 Constraint Networks of Topological Relations and Convexity. By E. Davis, N.M. Gotts and A.G. Cohn. CONSTRAINTS, Vol. 4 No. 3, 1999, pp. 241-280.
- J.10 Continuous Shape Transformations and Metrics on Regions, Fundamenta Informaticae, Vol. 46, Nos. 1-2, 2001, pp. 31-54.
- J.11 A First-Order Theory of Communication and Multi-Agent Plans. By E. Davis and L. Morgenstern. *Journal of Logic and Computation*, Vol. 15, No. 5, 2005, pp. 701-749. This paper has two online appendices, 7 and 16 pages long respectively, at http://cs.nyu.edu/faculty/davise/commplan-appa.pdf and commplan-appb.pdf.
- J.12 Knowledge and Communication: A First-Order Theory. Artificial Intelligence, vol. 166 nos. 1-2, 2005, pp. 81-140.
- J.13 Processes and Continuous Change in a SAT-Based Planner. By J. Shin and E. Davis. Artificial Intelligence, vol. 166 nos. 1-2, 2005, pp. 194-253.
- J.14 The Expressivity of Quantifying over Regions. Journal of Logic and Computation, vol. 16, 2006, pp. 891-916.
- J.15 Pouring Liquids: A Study in Commonsense Physical Reasoning. Artificial Intelligence, vol. 172, 2008, pp. 1540-1578. This paper has a 18 page online appendix at http://cs.nyu.edu/faculty/davise/papers/liqAppa.pdf.
- J.16 How Does a Box Work? A Study in the Qualitative Dynamics of Solid Objects. Artificial Intelligence, 175, 2011, pp. 299-345. This paper has a 20 page online appendix at http://cs.nyu.edu/faculty/davise/box-proof.pdf.
- J.17 Preserving Geometric Properties in Reconstructing Regions from Internal and Nearby Points. Computational Geometry: Theory and Applications, 45:5-6, 2012, 234-253.
- J.18 Elementarily Equivalent Domains for Topological Languages over Regions in Euclidean Space. Journal of Logic and Computation, 23:3, 2013, 457-471.
- J.19 Qualitative Spatial Reasoning in Interpreting Text and Narrative.
  Spatial Cognition and Computation, 13:4, 2013, 264-294.
  Also Space, Language, and Ontology: A Response to Bateman, same issue, 315-318.
- J.20 A Qualitative Calculus for Three-Dimensional Rotations. By A. Asl and E. Davis. Spatial Cognition and Computation, 14:1, 2014, 18-57.
- J.21 The Expressive Power of First-Order Topological Languages. Journal of Logic and Computation, 23:5, 2013, 1107-1141.
- J.22 How Robust Are Probabilistic Models of Higher-Level Cognition? By G. Marcus and E. Davis. *Psychological Science* 24:12, 2013, 2351-2360.

- J.23 The Singularity and the State of the Art in Artificial Intelligence, ACM Ubiquity, October 2014.
- J.24 The Scope and Limits of Simulation in Automated Reasoning. By E. Davis and G. Marcus. Artificial Intelligence, 233, April 2016, 60-72.
- J.25 Commonsense Reasoning and Commonsense Knowledge in Artificial Intelligence. By E. Davis and G. Marcus. Communications of the ACM, September 2015, 92-105.
- J.26 Still searching for principles: A response to Goodman et al. (2015). By G. Marcus and E. Davis. *Psychological Science*, 2015, 26:542-544.
- J.27 Lousy advice to the love-lorn. Communications of the ACM, December 2017.
- J.28 "Causal model" must be broadly construed. By E. Davis and G. Marcus. (Comment on B. Lake, T. Ullmann, J. Tenenbaum, and S. Gershman, "Building Machines that Learn and Think like People.") Behavioral and Brain Sciences, 40, 2017.
- J.29 Commonsense reasoning about containers using radically incomplete information. By E. Davis, G. Marcus, and N. Frazier-Logue. Artificial Intelligence, July 2017, 248, 46-84.
- J.30 Logical formalizations of commonsense reasoning: A survey. Journal of AI Research, August 2017, **59**, 651-723.
- J.31 Computational limits don't fully explain human cognitive limitations. by E. Davis and G. Marcus. (Comment on F. Lieder and T. Griffiths, "Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources") Behavioral and Brain Sciences, 2020, 43.
- J.32 Unanswerable questions about images and texts. Frontiers in Artificial Intelligence, July 2020.
- J.33 Broken Physics: A Conjunction-Fallacy Effect in Intuitive Physical Reasoning by Ethan Ludwin-Peery, Neil Bramley, Ernest Davis, and Todd Gureckis. *Psychological Sci*ence, November 2020.
- J.34 Insights for AI from the Human Mind. By G. Marcus and E. Davis. Communications of the ACM, January 2021, 61(1):38-41.
- J.35 Limits on Simulation Approaches in Intuitive Physics. by Ethan Ludwin-Peery, Neil Bramley, Ernest Davis, and Todd Gureckis. *Cognitive Psychology*, Vol. 127, June 2021.
- J.36 Benchmarks for Automated Commonsense Reasoning: A Survey. ACM Surveys, 56:4, Article 81, Feburary 2023, 41 pages.
- J.37 The Defeat of the Winograd Schema Challenge. By V. Kocijan, T. Lukasiewicz, E. Davis, G. Marcus, and L. Morgenstern. Artificial Intelligence, Vol. 325, December 2023.
- J.38 Mathematics, word problems, common sense, and artificial intelligence. Bulletin of the American Mathematical Society, Vol. 61 No. 2, April 2024, 287-303.

#### **Conference Proceedings**

- P.1 The MERCATOR Representation of Spatial Knowledge. Proceedings of the 8th IJCAI, 1983
- P.2 A Representation for Complex Physical Domains. By S. Addanki and E. Davis. Proceedings of the 9th IJCAI, pp. 443-446, 1985
- P.3 Inferring Ignorance from the Locality of Visual Perception. Proc. AAAI-88, pp. 786-790
- P.4 Error Correction in Cognitive Maps. Proc. Workshop on Sensor Fusion: Spatial Reasoning and Scene Interpretation, SPIE, 1988.
- P.5 Solutions to a Paradox of Perception with Limited Acuity. First International Conference on Knowledge Representation and Reasoning, 1989.
- P.6 The Semantics of Tasks that can be Interrupted or Abandoned. First International Conference on AI Planning Systems, 1992, pp. 37-44.
- P.7 Axiomatizing Qualitative Process Theory. Third International Conference on Knowledge Representation and Reasoning, 1992, pp. 177-188.
- P.8 Infinite Loops in Finite Time: Some Observations. Third International Conference on Knowledge Representation and Reasoning, 1992, pp. 47-58.
- P.9 Branching Continuous Time and the Semantics of Continuous Action. Second International Conference on AI Planning Systems, 1994, pp. 231-236.
- P.10 A First-Order Theory of Communicating First-Order Formulas. Ninth International Conference on Knowledge Representation and Reasoning, 2004. pp. 235-245. [Conference-length version of J.12]
- P.11 Continuous Time in a SAT-Based Planner. By J. Shin and E. Davis. Proc AAAI-2004. pp. 531-536. [Conference-length version of J.13]
- P.12 Ontologies and Representations of Matter. AAAI-10.
- P.13 The Winograd Schema Challenge. By H. Levesque, E. Davis, and L. Morgenstern. *KR-2012*.
- P.14 Reasoning from Radically Incomplete Information: The Case of Containers. By E. Davis, G. Marcus, and A. Chen. Advances in Cognitive Systems, 2013, 273-288.
- P.15 The Logical Depth of Reasoning about Other Minds. Advances in Cognitive Systems, 2017.
- P.16 Limits on the use of simulation in physical reasoning. By Ethan Ludwin-Peery, Neil Bramley, Ernest Davis, and Todd Gureckis. *Cognitive Science*. 2019.
- P.17 A generalization test of conjunction errors in physical reasoning. (Poster) By Ethan Ludwin-Peery, Neil Bramley, Ernest Davis, and Todd Gureckis. *Cognitive Science*. 2020.
- P.18 Physical reasoning in an open world. By Zhuoran (Jennifer) Zeng and Ernest Davis. Advances in Cognitive Systems, 2021.

#### P.19 My Experience Teaching Logic in Undergraduate AI at NYU. LogTeach-22. July 2022.

#### Chapters in Books

- C.1 Limits and Inadequacies in Artificial Intelligence. In No Way: On the Nature of the Impossible, Philip Davis and David Park, (eds.), W.H. Freeman, 1987, pp. 90-110
- C.2 Commonsense Reasoning. In *The Encyclopedia of Artificial Intelligence*, Stuart Shapiro, (ed.), John Wiley and Sons, 1987, pp. 833-840. Revised for second edition, 1990.
- C.3 A Framework for Qualitative Reasoning about Solid Objects. In J. de Kleer and D. Weld (eds.), *Readings in Qualitative Physical Reasoning*, Morgan Kaufmann, 1989, pp. 603-609. (Modified version of [J.5])
- C.4 Order of Magnitude Reasoning in Qualitative Differential Equations. In J. de Kleer and D. Weld (eds.), *Readings in Qualitative Physical Reasoning*, Morgan Kaufmann, 1989, pp. 422-434.
- C.5 Knowledge Representation. *The International Encyclopedia of the Social and Behavioral Sciences*, Neil J. Smelser and Paul B. Baltes (eds), Elsevier Science Pubs, Oxford, 2001, pp. 8132-8139. Revised for 2nd edition, 2015, Vol. 13, pp. 98-104.
- C.6 Physical Reasoning. The Handbook of Knowledge Representation, F. van Harmelen, V. Lifschitz, and B. Porter (eds.), Elsevier, Oxford, 2008, chap. 14, pp. 597-620.
- C.7 Qualitative Reasoning and Spatio-Temporal Continuity. *Qualitative Spatio-Temporal Representation and Reasoning: Trends and Future Directions*, S. Hazarika (ed.), IGI Global, 2012.
- C.8 Introduction. In Mathematics, Substance and Surmise: Views on the Meaning and Ontology of Mathematics, E. Davis and P. Davis (editors), Springer, 2015.
- C.9 How Should Robots Think about Space? In *Mathematics, Substance and Surmise: Views on the Meaning and Ontology of Mathematics, E. Davis and P. Davis (editors), Springer, 2015.*
- C.10 Proof Verification Technology and Elementary Physics. In Algorithms and Complexity in Mathematics, Epistemology, and Science, N. Fillion, R. Corless, and I.S. Kotsireas (eds.) Springer, 2019.
- C.11 Using human skills taxonomies and tests as measures of AI. In AI and the Future of Skills, Vol. 1: Capabilities and Assessments. ed. Stuart Elliott, OECD Publishing, 2021.

#### Other contributions to books

D.1 47 original exercises with solutions in S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach* 3rd edn. and accompanying instructor's manual, Prentice Hall, 2009.

#### Reviews

- R.1 Response to "Prolegomena to any future Qualitative Physics," by Elisha Sacks and Jon Doyle. Computational Intelligence, vol. 8, no. 2, 1991, pp. 316-318.
- R.2 Review of *Mirror Worlds*, by David Gelernter, Oxford U. Press, 1991. SIAM News, vol. 25, no. 3, 1992, p. 6.
- R.3 Response to reviews of Representations of Commonsense Knowledge. Artificial Intelligence, vol. 61, 1993, pp. 175-179.
- R.4 Review of Qualitative Reasoning, by Benjamin Kuipers, MIT Press. IEEE Expert, vol. 9, no. 6, 1994, pp. 70-71.
- R.5 Review of The Uncertain Reasoner's Companion: A Mathematical Perspective, by J.B. Paris, Cambridge University Press. IEEE Expert, vol. 10, no. 5, 1995, pp. 78-79.
- R.6 Gödel, Escher, Bach redux. IEEE Expert, vol. 11, no. 3, 1996, pp. 7, 9.
- R.7 Two Machine Learning Textbooks: An Instructor's Perspective (Reviews of Machine Learning by Tom Mitchell and of Data Mining by Ian H. Witten and Eibe Frank). Artificial Intelligence, vol. 131, Sept. 2001, pp. 191-198.
- R.8 Mathematics as Metaphor: Review of Where Mathematics Comes From, by George Lakoff and Raphael Nuñez. Journal of Experimental and Theoretical Artificial Intelligence, vol. 17, no. 3, 2005, pp. 305-315.
- R.9 Review of In the Land of Invented Languages, by Arika Okrent, SIAM News, June, 2010, pp. 4,6.
- R.10 Review of Donald Michie on Machine Intelligence, Biology, and More, ed. Ashwin Srinivasan, Times Literary Supplement, July 16, 2010, p. 27.
- R.11 Review of *The Invisible Gorilla: And Other Ways our Intuitions Deceive Us*, by Christopher Chabris and Daniel Simons, *American Scientist*, Sept-Oct. 2010, **98**:5, p. 428-430.
- R.12 Review of Roads to Infinity: The Mathematics of Truth and Proof by John Stillwell. SIAM News, Vol. 43, No. 8, October 2010, pp. 6,7.
- R.13 Review of Galileo, by J.L. Heilbron, SIAM News, 44:4, May 2011.
- R.14 Review of *The Information: A history, a theory, a flood, by James Gleick, Times Literary Supplement, August 17, 2011, #5655-6, p. 10.*
- R.15 Review of Geometry Revealed: A Jacob's Ladder to Modern Higher Geometry by Marcel Berger, SIAM News, 44:7, September 2011.
- R.16 Review of Metareasoning: Thinking about Thinking ed. Michael Cox and Anita Raja, Computing Reviews, August 12, 2011, Review #139347.
- R.17 Review of "A note on a generalization of the muddy children problem" by N. Gierasimczuk and J. Szymanik, *TARK 13*, 2011. Review #139441 in *Computing Reviews*, Sept. 9, 2011.
- R.18 Review of Language technology for cultural heritage : selected papers from the LaTeCH Workshop Series ed. C. Sporleder, Antal van den Bosch and K. Zervanou. Computing Reviews, Nov. 2, 2011, Review #139550.
- R.19 Review of Nine Algorithms that Changed the Future: The Ingenious Ideas that Drive Today's Computers, by John MacCormick, SIAM News, March, 2012.

- R.20 Review of "Robust grasping under object pose uncertainty," by Kaijen Hsaio, Leslie Pack Kaelbling, and Tomas Lozano-Pérez, Automonomous Robots, 2011, 31:253-268. Review #CR140061 in Computing Reviews, April 2012.
- R.21 Review of In Pursuit of the Travelling Salesman: Mathematics at the Limits of Computation, by William Cook. SIAM News, May 2012.
- R.22 Review of *Chasing Venus: The Race to Measure the Heavens*, by Andrea Wulf. *SIAM News*, July/August 2012.
- R.23 Review of Computation and its Limits by Paul Cockshott, Lewis Mackenzie and Greg Michaelson. SIAM News, November 2012.
- R.24 Review of "Exploiting Qualitative Spatial Reasoning for Topological Adjustment of Spatial Data," by Jan Oliver Wallgrün, ACM SIGSPATIAL GIS 2012. Review #CR140872 in Computing Reviews. January, 2013.
- R.25 "What Nate Silver Gets Wrong" by G. Marcus and E. Davis. Review of *The Noise and the Signal*, by Nate Silver. Review in the *New Yorker online edition*, Jan. 25, 2013.
- R.26 Review of Special Issue of Artificial Intelligence Journal on AI and Wikipedia, Vol. 194 No. 1. Edited by E. Hovy, R. Navigli, and S.P. Ponzetto. Review #CR140920, Computing Reviews, Feb. 8, 2013.
- R.27 Review of Newton and the Origin of Civilization, by Jed Z. Buchwald and Mordechai Feingold. SIAM News, May 1, 2013.
- R.28 Review of Jane Austen: Game Theorist, by Michael Chwe. SIAM News, September 2013.
- R.29 Review of Probably Approximately Correct, by Leslie Valiant. SIAM News, October 2013.
- R.30 Review of Case-based reasoning: A textbook, by M. Richter and R. Weber. Review #CR141932, Computing Reviews, Jan. 27, 2014.
- R.31 Review of Who's Bigger? Where Historical Figures Really Rank, by S. Skiena and C. Ward. SIAM News, March 2014.
- R.32 Review of Three views of logic: Mathematics, philosophy, and computer science, by D. Loveland, R. Hodel, and S.G. Sterrett. Review #CR142187, Computing Reviews, April 17, 2014.
- R.33 Review of "An Empirical Perspective on Representing Time" by A. Scheuermann et al. Review #CR142757, *Computing Reviews*, August 1, 2014.
- R.34 Review of How Not to be Wrong, by Jordan Ellenberg. SIAM News, September 2014.
- R.35 Ethical Guidelines for a Superintelligence. Review of SuperIntelligence: Paths, Dangers, Strategies, by Nick Bostrom. Artificial Intelligence 220, January 2015, 121-124.
- R.36 Review of Mathematics without Apologies: Portrait of a Problematic Vocation, by Michael Harris. SIAM News, March 2015.
- R.37 Review of Numerical Notation: A Comparative History, by Stephen Chrisomalis. SIAM News, June 2015.
- R.38 Review of *Machine Translation*, by P. Bhattacharya. Review #CR143722, *Computing Reviews*, August 2015.

- R.39 Review of Through the Language Glass: Why the World Looks Different in Different Languages, by Guy Deutscher and The Language Hoax: Why The World Looks the Same in Any Language, by John McWhorter. Artificial Intelligence, December 2015, 229:202-209.
- R.40 Review of The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World, by Pedro Domingos. SIAM News, January 2016.
- R.41 Review of Dance Notations and Robot Motions ed. Jean-Paul Laumond and Naoko Abe, Review #CR144189, Computing Reviews, February 2016.
- R.42 Review of Algorithms to Live By, by Brian Christian and Tom Griffiths, Artificial Intelligence, 239, 1-6, October 2016.
- R.43 Review of Nathaniel Bowditch and the Power of Numbers, by Tamara Plakins Thornton. SIAM News, July 2016.
- R.44 Review of Leonhard Euler: Mathematical Genius in the Enlightenment, by R. Callinger. Review #CR144743, Computing Reviews, September 2016.
- R.45 Review of ENIAC in Action: Making and Remaking the Modern Computer, by Thomas Haigh, Mark Priestley, and Crispin Rope. SIAM News, October 2016.
- R.46 Review of *The Turing Guide*, by Jack Copeland, Jonathan Bowen, Mark Sprevak, and Robin Wilson. *SIAM News*, July/August 2017. Chinese translation published in *Mathematics*, 37:1.
- R.47 Review of A Mind at Play: How Claude Shannon Invented the Information Age, by Jimmy Soni and Rob Goodman. SIAM News, November 2017.
- R.48 Review of *Roman Portable Sundials: The Empire in Your Hand*, by Richard J.A. Talbert. Unpublished. January 2018.
- R.49 Review of Exact Thinking in Demented Times: The Vienna Circle and the Epic Quest for the Foundations of Science, by Karl Sigmund. SIAM News. February 2018.
- R.50 Review of Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor, by Virginia Eubanks. SIAM News, May 2018.
- R.51 Review of The Pendulum Paradigm: Variations on a Theme and the Measure of Heaven and Earth by Martin Beech. SIAM News, April 2019.
- R.52 Review of The Life and Work of George Boole: A Prelude to the Digital Age, by Desmond MacHale; New Light on George Boole, by Desmond MacHale and Yvonne Cohen; and The Continued Exercise of Reason: Public Addresses by George Boole, ed. Brendan Dooley. SIAM News, June 2019.
- R.53 Review of Laura Bassi and Science in 18th Century Europe, by Monique Frize. SIAM News, September, 2019.
- R.54 Review of *How to Walk on Water and Climb Up Walls: Animal Movements and the Robots of the Future*, by David Hu. SIAM News, November 2019.
- R.55 Review of Shape: The Hidden Geometry of Information, Biology, Strategy, Democracy, and Everything Else, by Jordan Ellenberg. SIAM News, July 2021.
- R.56 Review of Journey to the Edge of Reason: The Life of Kurt Gödel, by Stephen Budiansky. SIAM News, October 2021.

- R.57 Review of Curves for the Mathematically Curious: An Anthology of the Unpredictable, Historical, Beautiful, and Romantic by Julian Havil. SIAM News, December 2021.
- R.58 Review of Uncountable: A Philosophical History of Number and Humanity from Antiquity to the Present by David Nirenberg and Ricardo Nirenberg. SIAM News, March 2022.
- R.59 Review of Mathematician with the Soul of a Poet: Poems and Plays of Sofia Kovalevskaya. translated, edited, and introduced by Sandra DeLozier Coleman. SIAM News, April 2022.
- R.60 Review of *Rules: A Short History of What We Live By* by Lorraine Daston. *SIAM News*, October 2022.
- R.61 Review of John Venn: A Life in Logic by Lucas Verburgt. SIAM News, November 2022.
- R.62 Review of The Story of Proof: Logic and the History of Mathematics, by John Stillwell. SIAM News, January 2023.
- R.63 Review of Axiomatics: Mathematical Thought and High Modernism by Alma Steingart. SIAM News, March 2023.
- R.64 Review of *Phenomena: Doppelmayr's Celestial Atlas* by Giles Sparrow. *SIAM News*, April 2023.
- R.65 Review of Data Driven: Truckers, Technology, and the New Workplace Surveillance, by Karen Levy. SIAM News, October 2023.
- R.66 Review of Unmasking AI: My Mission to Protect What is Human in a World of Machines by Joy Buolamwini, and Your Face Belongs to Us: A Secretive Startups Quest to End Privacy as We Know It by Kashmir Hill. SIAM News, March 2024.
- R.67 Review of The Worlds I See: Curiosity, Exploration, and Discovery at the Dawn of AI by Fei-Fei Li. SIAM News, May 2024.

#### Published Articles without Peer Review

- A.1 AI Research in Progress at the Courant Institute, New York University. Edited by E. Davis and R. Grishman. *AI Magazine*, Winter, 1986, Vol. 7, No. 5, pp. 82-86
- A.2 Qualitative simulation and prediction. *IEEE Expert*, vol. 12, no. 3, 1997, p. 103.
- A.3 The Naive Physics Perplex. AI Magazine, Winter 1998, Vol. 19. No. 4. pp. 51-79.
- A.4 Guide to Axiomatizing Domains in First Order Logic E. Sandewall (ed.) *Electronic Newsletter on Reasoning about Actions and Change*, Issue 99002, March 8, 1999.
- A.5 Progress in Formal Commonsense Reasoning. By E. Davis and L. Morgenstern. Artificial Intelligence, vol. 153, March 2004, nos. 1-2, pp. 1-12.
- A.6 Reports of the AAAI 2011 Spring Symposia. M. Buller et al. (19 authors, including E. Davis). AI Magazine, Fall 2011, Vol. 32, No. 3, 119-127.

- A.7 Expand Minds, Not the NYU Campus.
  - By E. Davis, P. Deer, and M.C. Miller. New York Times Op-Ed, April 26, 2012. Reprinted in While We Were Sleeping: NYU and the Destruction of New York, ed. M.C. Miller, McNally-Jackson, 2012.
- A.8 Maths is the true language of science.By G. Marcus and E. Davis. *Financial Times*, April 13, 2013.
- A.9 A Grand Unified Theory of Everything.By G. Marcus and E. Davis, New Yorker, online edition May 6, 2013.
- A.10 Eight (No, Nine!) Things Wrong with Big Data. By G. Marcus and E. Davis, New York Times, Op-Ed, April 7, 2014.
- A.11 Do We Really Need to Learn to Code? By G. Marcus and E. Davis, New Yorker, online edition June 7, 2014.
- A.12 How to Write Science Questions that are Easy for People and Hard for Computers. AI Magazine, Spring 2016, 13-22.
- A.13 Planning, Executing, and Evaluating the Winograd Schema Challenge. By L. Morgenstern, E. Davis, and C. Ortiz. AI Magazine, Spring 2016, 50-54.
- A.14 The Tragic Tale of Tay the Chatbot. AI Matters, 2:4, Summer 2016.
- A.15 Managing Big Data's Big Risks. *Project Syndicate*, Feb. 2017.
- A.16 The First Winograd Schema Challenge at IJCAI-16. By E. Davis, L. Morgenstern, and C. Ortiz. AI Magaine, Fall 2017, 97-98.
- A.17 AI is harder than you think. By G. Marcus and E. Davis. New York Times, May 2017.
- A.18 No, A.I. won't solve the fake news problem. By G. Marcus and E. Davis. New York Times, October 2018.
- A.19 How to build artificial intelligence we can trust. By G. Marcus and E. Davis. New York Times, September 2019.
- A.20 If computers are so smart, how come they can't read? By G. Marcus and E. Davis. WIRED September 2019.
- A.21 Six questions to ask yourself when reading about AI. By G. Marcus and E. Davis. *Quartz*, September 2019.
- A.22 AI's hardest problem? Developing common sense. By G. Marcus and E. Davis. *LinkedIn*, October 2019.
- A.23 Are Neural Networks About to Reinvent Physics? By G. Marcus and E. Davis. *Nautilus*, November 2019.
- A.24 Q&A on the Book Rebooting AI. By G. Marcus and E. Davis. *InfoQ*, November 2019.
- A.25 GPT-3, Bloviator: OpenAI's language generator has no idea what it's talking about. By G. Marcus and E. Davis, *Technology Review*, August 2020.

- A.26 Has AI Found a New Foundation?By G. Marcus and E. Davis, *The Gradient*, September 2021.
- A.27 What does Meta AI's Diplomacy-winning Cicero Mean for AI? by G. Marcus and E. Davis. Comm. ACM blog, November 2022.
- A.28 Large Language Models like ChatGPT say the Darnedest Things. by G. Marcus and E. Davis. Comm. ACM blog, January 2022.
- A.29 How Not To Test GPT-3. By G. Marcus and E. Davis. Comm. ACM blog, February 2023.

# Technical reports, workshop notes, arXiv articles, and other unpublished material, not otherwise listed

- T.1 Organizing Spatial Knowledge Tech. Report #193, Yale Computer Science Dept., January 1981
- T.2 An Ontology of Physical Action Tech. Rep. 123, NYU Comp. Sci. Dept., June 1984
- T.3 Shape and Function of Solid Objects: Some Examples Tech. Rep 137, NYU Comp. Sci. Dept., October 1984
- T.4 A High Level Real-Time Programming Language Tech. Rep. 145, NYU Comp. Sci. Dept., October 1984
- T.5 Reasoning about Hand-Eye Coordination. IJCAI-89 Workshop on Knowledge, Perception, and Planning
- T.6 Physical Idealization as Plausible Inference.
  Tech. Rep. 534, NYU Comp. Sci. Dept., December, 1990.
  Logical Formalisms of Commonsense Reasoning, Stanford Spring Symposium, 1991.
- T.7 Lucid Representations. Tech. Rep. 565, NYU Comp. Sci. Dept., June 1991.
- T.8 Epistemic Logic and Its Applications (with L. Morgenstern). Epistemic Logic: Annotated Bibliography. Tutorial Notes, IJCAI-93.
- T.9 Approximations of Shape and Configuration Space. NYU Computer Science Tech. Report #703, September 1995. Extensively revised under the title "Kinematic Tolerance and the Topology of Configuration Space," at http://www.cs.nyu.edu/faculty/davise/papers/path.pdf, 2007.
- T.10 Approximation and Abstraction in Solid Object Kinematics. NYU Computer Science Tech. Report #706, September 1995.
- T.11 A Highly Expressive Language of Spatial Constraints. NYU Computer Science Tech. Report #714, December 1995.
- T.12 Describing spatial transitions using mereotopological relations over histories. NYU Computer Science Tech. Report #2000-809, October 2000.
- T.13 Collection of Winograd Schemas. At http://www.cs.nyu.edu/faculty/davise/papers/WS.html. Created September 2011 and updated since.

- T.14 The Relevance of Proofs of the Rationality of Probability Theory to Automated Reasoning and Cognitive Models. Unpublished. arXiv:1310.1328
- T.15 Recent Critiques of Big Data: Bibliography. http://www.cs.nyu.edu/faculty/davise/papers/BigDataBib.html
- T.16 Contempt for conservatism, religion, and history among social scientists. Guest blog in Lee Jussim's blog "Rabble Rouser;" in *Psychology Today*, November 2014.
- T.17 The Limitations of Standardized Science Tests as Benchmarks for Artificial Intelligence Research: Position Paper. arXiv:1411.1629. November, 2014.
- T.18 The Scope and Limits of Simulation in Cognitive Models. By E. Davis and G. Marcus. arXiv:1506.04956
- T.20 Human tests of materials for the Winograd Schema Challenge. By E. Davis, L. Morgenstern, and C. Ortiz. Unpublished.
- T.21 Winograd Schemas and Machine Translation. arXiv:1608.01884. August 2016.
- T.22 Google Translate fails on simple sentences. Initially created October 2016, and intermittently updated.
- T.23 In Praise of Globes. Guest blog in Cathy O'Neil's blog "mathbabe". March 23, 2017.
- T.24 On Deep Learning for Symbolic Integration: A review of (Lample and Charton, 2019). arXiv:1912.05752. December 2019.
- T.25 A Review of Winograd Schema Challenge Databases and Approaches. By V. Kocijan, T. Lukasiewicz, E. Davis, G. Marcus, and L. Morgenstern. arXiv:2004.13831. April 2020.
- T.26 The test set for the TransCoder System. arXiv:2008.00293. August 2020.
- T.27 A flawed dataset for symbolic equation verification. arXIv:2105.11479. May 2021.
- T.28 Metric topologies over some categories of simple open regions in Euclidean space. arXiv:2109.08663. September 2021.
- T.29 Deep Learning and Mathematical Intuition: A review of (Davies et al. 2021). arXiv:2112.04324. December 2021.
- T.31 Pragmatic constraints and pronoun reference disambiguation: the possible and the impossible. arXiv:2204.01166. April 2022.
- T.32 A very preliminary analysis of DALL-E 2. By G. Marcus, E. Davis, and S. Aaronson. arXiv:2204.13807. May 2022.
- T.34 Experiments in Commonsense Reasoning in GPT-3: Status Report from June 2022. By E. Davis and G. Marcus. June 2022.
- T.35 What has surprised me in computer technology. June 2022.
- T.36 "Learning by heart" in many languages. August 2022.
- T.37 Limits of an AI program for solving college math problems. arXiv:2208.06906 August 2022.
- T.38 Testing GPT-4 with Wolfram Alpha and Code Interpreter plug-ins on math and science problems. By E. Davis and S. Aaronson. arXiv 2308.05713. August 2023.

- T.39 Reports of the birth of AGI are greatly exaggerated. By G. Marcus and E. Davis. Substack blog. October 2023.
- T.40 Hello, Multimodal Halluciations By G. Marcus and E. Davis. Substack blog. October 2023.
- T.41 Getting GPT to work with external tools is harder than you think. By G. Marcus and E. Davis. Substack blog. October 2023.
- T.42 Comment on (Romera-Paredes et al., 2023), "Mathematical discoveries from program search with large language models".