Curriculum Vitae

Adam Kalai

January, 2001

Work:
Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213

Home:
4902 Forbes Ave. #195
Pittsburgh, PA 15213

cell: (412) 901-5202
fax: (412) 268-5576
akalai@cs.cmu.edu
http://www.cs.cmu.edu/~akalai

Research Interests

Algorithms: on-line, randomized.
Machine learning: computational learning theory, web search, data mining.
3D autostereoscopic displays.

Education

Harvard University — Cambridge, MA

B.A. Computer Science, 1996, Magna Cum Laude.

Carnegie Mellon University — Pittsburgh, PA

M.A. Computer Science, 1999.
Ph.D. Computer Science, expected May 2001.
Advisor: Avrim Blum.
Thesis title: "Probabilistic and On-line Methods in Machine Learning."
Thesis committee: Avrim Blum, Manuel Blum, Danny Sleator, and Santosh Vempala.
Support: NSF Graduate Fellowship and IBM Distinguished Graduate Fellowship.

Honors and Awards

IBM Distinguished Graduate Fellowship, 1999-2001.
NSF Graduate Fellowship, 1996-1999.
Second Place, ACM World Programming Contest, 1996 (3-person Harvard team).
First Place, Harvard Programming Contest, 1994.
Eighteenth Place, William Lowell Putnam Math Contest 1994.
First Place, Illinois Junior Invitational Chess Championship, 1992.

Work Experience

9/1996–5/2001 Graduate Student, Carnegie Mellon University — Pittsburgh, PA

Research in machine learning theory and on-line algorithms with Avrim Blum. Other projects include 3D reconstruction with Steve Seitz, 3D displays with Mel Siegel, cryptography with Manuel Blum, and virtual reality with Randy Pausch.

6/2000-7/2000 Summer Intern, IBM Almaden — San Jose, CA

Design and analysis of efficient algorithms for fair scheduling, specifically the "car pool problem."

5/1999–8/1999 Summer Intern, IBM Watson — Hawthorne, NY

Design, analysis, and implementation of automatic pricing strategies for so-called information economies.

6/1998–9/1998 Summer Intern, Xerox PARC — Palo Alto, CA

Sole design and implementation of an intelligent web search engine (like CLEVER) from the raw data of over fifty million web pages and three hundred gigabytes of HTML.

2/1996–6/1996 Database Programming, Northwestern University — Evanston, IL

Relational database design and programming for the journal *Games and Economic Behavior.*

9/1995–1/1996 Programmer, Harvard University — Cambridge, MA

Team design and implementation of a system on efficient distributed execution of programs, under Professor Michael Rabin.

6/1995–8/1995 Programmer, Kurzweil Applied Intelligence — Waltham, MA

Design and implementation of a natural language processing section of a speech recognizer.

Teaching Experience

Fall 2000 Substitute Teaching, Carnegie Mellon University — Pittsburgh, PA

Occasional substitute teaching for Steven Rudich's graduate "Complexity Theory."

Spring 1998 Teaching Assistant, Carnegie Mellon University — Pittsburgh, PA

Teaching assistant for undergraduate "Mathematical Foundations of Computer Science," a course by Steven Rudich and Bruce Maggs. This included teaching a weekly one-hour section.

Fall 1997 Teaching Assistant, Carnegie Mellon University — Pittsburgh, PA

Teaching assistant for Danny Sleator's undergraduate "Algorithms."

Fall 1995 Teaching Fellow, Harvard University — Cambridge, MA

Teaching fellow for Michael Rabin's graduate "Efficient Algorithms."
(A teaching assistant at Harvard is a teaching fellow.)

Summers 1992, 1993 Teaching Assistant, Center for Talent Development — Evanston, IL

Teaching assistant for talented junior high school students studying algebra and geometry.

Other Professional Activities

Referee for papers in the following journals and conferences:

Machine Learning, *IEEE Transactions on Information Theory*, *Information Processing Letters*, *Journal of Algorithms*, and *Journal of Computer and System Sciences*, *Journal of Artificial Intelligence Research*, AsiaCrypt 2000, FOCS 2000.

Tetric – the 2000 IC programming contest.

Co-organized with Danny Sleator a fun contest in which teams wrote programs that play tetris. Goal was to draw pictures by playing tetris. Details and Java applet at http://www.cs.cmu.edu/~akalai/tetric.

Publications

[KV00] Adam Kalai and Santosh Vempala, "Efficient Algorithms for Universal Portfolios," Proceedings of the 41st Annual Symposium on the Foundations of Computer Science (FOCS '00), Redondo Beach, 2000.

[KK00] Adam Kalai and Ehud Kalai, "Strategic Polarization," to appear in *Journal of Mathematical Psychology*.

[BKW00] Avrim Blum, Adam Kalai, and Hal Wasserman, "Noise-Tolerant Learning, the Parity Problem, and the Statistical Query Model," Proceedings of the 32nd Annual Symposium on the Theory of Computing (STOC '00).

[K00] Adam Kalai, "Generating Factored Numbers, Easily," submitted to *Journal of Cryptology*.

[BKL99] Avrim Blum, Adam Kalai, and John Langford, "Beating the Holdout: Bounds for K-Fold and Progressive Cross-Validation," Proceedings of the Tenth Annual Conference on Computational Learning Theory (COLT '99).

[BBK99] Avrim Blum, Carl Burch, and Adam Kalai, "Finely Competitive Paging," Proceedings of the 40th Annual Symposium on the Foundations of Computer Science (FOCS '99).

[BK99] Avrim Blum and Adam Kalai, "Universal Portfolios With and Without Transaction Costs," *Machine Learning*, 35:3 pp 193-205, 1999.

[SKS99] Harry Shum, Adam Kalai, and Steve Seitz, "Omnivergent Stereo," Proceedings of the Sixth International Conference on Computer Vision (ICCV '99).

[KCBR99] Adam Kalai, Stanley Chen, Avrim Blum, and Roni Rosenfeld, "On-line Algorithms for Combining Language Models," Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP '99).

[KS98] Adam Kalai and Mel Siegel, "Improved Rendering of Parallax Panoramagrams for a Time-Multiplexed Autostereoscopic Display," Proceedings of SPIE 3295, Stereoscopic Displays and Applications IX, 1998.

[BK98] Avrim Blum and Adam Kalai, "A Note on Learning from Multiple-Instance Examples," *Machine Learning*, 30:1, pp 23-30, 1998.

Publication Summaries

[KV00] "Efficient Algorithms for Universal Portfolios"

The Universal Portfolio is an investment strategy introduced by Thomas Cover. Previously, all known implementations required computation exponential in the number of stocks. In this paper, we use random walks for an implementation that is polynomial in the number of stocks.

[KK00] "Strategic Polarization"

An example of polarization is a marriage in which one spouse gives generously to charity while the other donates nothing. This may misrepresent what is, in actuality, a small discrepancy in preferences. It may be that the donating spouse would like to see

10% of their combined income go to charity each year, while the apparently frugal spouse would like to see 9% donated. A simple game-theoretic analysis suggests that the spouses will end up donating 10% and 0%, respectively. This paper gives a strat

egic (game-theoretic) analysis of polarization.

[BKW99] "Noise-Tolerant Learning, the Parity Problem, and the Statistical Query Model"

The "Statistical Query" model of Kearns has been essentially the only tool for noisy learning in computational learning theory. We present the first known problem that can be learned in the presence of noise but provably not by Statistical Queries.

[K00] "Generating Factored Numbers, Easily"

Our goal is to generate a uniformly random number between 1 and N, along with its prime factorization. Since there are no known efficient factoring algorithms, one cannot simply generate a random number and then factor it. Eric Bach's award-winning dissertation introduces and solves this problem. We present a four-line solution that is still efficient.

[BKL99] "Beating the Holdout: Bounds for K-Fold and Progressive Cross-Validation"

K-Fold cross-validation is a popular technique in machine learning for estimating the performance of a learned hypothesis on a data set. We provide the first theoretical justification for this method, showing that it is, on average, more accurate than a held-out test set of comparable size.

[BBK99] "Finely Competitive Paging"

In the on-line paging problem, one has a fixed-size cache. Each time a page is requested that is not in the cache, it has to be brought in at a cost of 1 and another page has to be evicted from the cache. We present an algorithm that has the same worst-case performance as the previous algorithms, but for a large and important class of request sequences, is superior.

[BK99] "Universal Portfolios With and Without Transaction Costs"

We present a much simpler analysis of Thomas Cover's Universal Portfolios. This analysis easily extends to the case of transaction costs, and eventually leads to "Efficient Algorithms for Universal Portfolios" (see above).

[SKS99] "Omnivergent Stereo"

We describe the design of an optimal camera for 3D reconstruction. In our model, a camera may capture a fixed number of pixels (light rays) from anywhere inside the circular frame of the camera. Afterwards, points outside the camera must be reconstructed as accurately as possible. We show that capturing tangent rays (in both directions) gives optimal uniform reconstructions.

[KCBR99] "On-Line Algorithms for Combining Language Models"

We show how Cover's Universal Portfolios can be applied to the field of language modeling. The goal of a language model is to accurately predict the next word in a sequence of words, with applications to speech recognition and data compression. In this setting, the universal guarantees are very striking. We support this with experiments on several different corpora.

[KS98] "Improved Rendering of Parallax Panoramagrams for a Time-Multiplexed Autostereoscopic Display"

We describe a new kind of 3D display that is viewable without glasses.

[BK98] "A Note on Learning from Multiple-Instance Examples"

In the multiple-instance learning setting introduced by Dietterich et al, the example data consist of n-tuples of points. An n-tuple is labeled positive if any of the n constituent points are positive. We show that this problem can be reduced to that of learning in the presence of noise. Thus, anything learnable in the presence of noise (including everything previously known to be multiple-instance learnable) is multiple-instance learnable.

References

Professor Avrim Blum
Department of Computer Science
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
(412) 268-6452
avrim@cs.cmu.edu
Exec Asst: daz@cs.cmu.edu

Professor Manuel Blum
Department of Computer Science
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
(412) 268-3742
mblum@cs.cmu.edu
Exec Asst: plp@cs.cmu.edu

Professor Steven Rudich
Department of Computer Science
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
(412) 268-7885
rudich@cs.cmu.edu
Exec Asst: daz@cs.cmu.edu

Professor Santosh Vempala
Deptartment of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139
(617) 253-4064
vempala@math.mit.edu