Greg P. Kochanski

48 Westminster Way,
Oxford OX2 0LW
England

gpk@kochanski.org

Specialities

◇

◇ Machine learning systems and statistical analyses.

◇ Many-gigabyte data sets; data mining.

◇ Software Development (Python, C++, other).

◇ Mathematical modeling of complex systems, including human behavior and language.

◇ Technical management.

Employment

◇ **University of Oxford**, 6/2003-present Oxford, UK
Senior Research Fellow, Oxford University Phonetics Laboratory; speech and language research.
Skills and Technologies:

- Natural language processing and language models.
- Signal and image processing.
- Building and using machine learning systems. Hidden Markov models.
- Linux scripting and programming; computational experiments.
- High performance computing.
- Statistics, especially Bayesian Markov-chain Monte Carlo.
- Mathematical models of the strategies humans use to produce speech.
- Project planning, management, experimental design, data analysis.

Major projects: (Details at `http://kochanski.org/gpk/kochanski_cv.pdf` and `http://kochanski.org/gpk/papers`.)

- How do we make compromises when we push words together? (Co-Investigator) What do you do if one word ends with the tongue low and the next starts with the tongue high?
- Techniques for effective measurement of speech rhythm. (Principal Investigator) How can you objectively define rhythmic speech?
- Mining a Year of Speech (Co-Investigator) How well can we find and display words within a large (terabyte) speech corpus?
- How do we construct the pronunciations of a word we've never said before? (Principal Investigator) Do we use rules or do we borrow pieces of previous words?
- Using dynamic MRI imaging to watch the tongue move and modeling the control of tongue position. (Principal Investigator)
- Applications for audio search tools in the Humanities. (Co-investigator)
- Consulting work on laser-based micro projectors. (Samsung).

- Human memories of intonation contours. (Principal Investigator) How much information does our brain retain?
- Models of English intonation. Behaviors used for expressing emphasis.

◇ **Rutgers University**, 9/2002–6/2003 New Brunswick, NJ
Visiting scholar, department of Chemistry; also consulting scientist for UCSD Dept. of Mechanical and Aerospace Engineering.

◇ **Bell Laboratories/Lucent Technologies**, 5/1987–9/2002 Murray Hill, NJ
Speech technology, computational physics, and experimental physics. As a Member of Technical Staff, contributed to the following projects:

- **Speech Synthesis R&D**, 1996–9/2002: Modeling of speech and intonation for the Bell Labs speech synthesis system. We used novel physiologically-based models to produce the world's best models of Chinese tones. Speech synthesizer design and data collection. Developed acoustic reverse Turing test (CAPTCHA). Microphone arrays and signal processing for improved speech data acquisition: echo suppression and noninvasive measurements of the larynx. Natural language processing.
- **Astrophysics**, 1994–1999: Observational astrophysics, using strong gravitational lensing as a probe of the mass distribution in clusters of galaxies. (Image processing & nonlinear optimization.)
- **Classifiers and Human Interfaces**, 1995–1996: Human interface studies and machine learning systems. Stylus keyboard design. Fruit recognition via bar-code scanner.
- **Nanomaterials and Condensed Matter Physics**, 1987–1999: R&D on flat-panel field-emission displays. Led team of 6 researchers. Measurements of material properties: Buckyballs and crystal growth in ultra-high vacuum. Designed, built, and used an early Scanning Tunneling Microscope, including software for real-time control of tip position with sub-nanometer precision.

Education

◇ **Massachusetts Institute of Technology**, Cambridge, MA
Ph.D. Physics 1987: "Magnetic Trapping and Cooling of Atomic Hydrogen". Built a novel low-temperature hydrogen maser; mathematical models of the momentum and spin distributions.

◇ **Massachusetts Institute of Technology**, Cambridge, MA
S.B. Physics 1982. First place, Mechanical Engineering Design Contest. Computer science and Engineering courses.

Other

◇ Attracted well over £1M of research funding to Oxford.

◇ Developed and taught courses and tutorials at University of Oxford (2004-):
∘ Statistics; ∘ Psycholinguistics; ∘ General Linguistics.

◇ Developed course and taught at Montclair State University (2002): Natural language processing. Taught at two international workshops (2003, 2011).

◇ Supervising two D. Phil (Ph.D.) students at University of Oxford.

◇ Montclair State University (NJ) Linguistics Department Advisory Board, 2002, 2003.

◇ Chair of two and member of 6 appointment committees

◇ Reviewed research grants for UK and French funding agencies; reviewed papers for at least 6 journals.

Skills.

◇ Python, Cython, C++, C, Bash, Java, Awk, R, FORTRAN, Limbo, LabView, Javascript, Perl.

◇ 10 years experience in Python: 200 kloc, of which 90 k are unique lines and and 60 kloc are open-sourced.

◇ XML, SQL, git (git hooks for automated testing), bzr, hg, svn, LAPACK, PORT, fftw, numpy, MPI, OpenMP, small network sysadmin.

◇ Large scale numeric computations; parallel code; threads. Condor and Beowulf clusters.

◇ Pilot (500 hours flight time); rock climbing; community theatre.

Selected Publications

◇ Thirty-five peer-reviewed technical papers, plus conference papers and talks. At least 2700 papers refer to mine.

"Prosodic peak estimation under segmental perturbations", Greg Kochanski *J. Acoust. Soc. Am.* 2010, 127(2), p. 862–873 doi:10.1121/1.3268511

Measuring Phonetic Distance: An example of Monte Carlo Optimization 5 May 2011, at the University of Illinois Speech Production Workshop, University of Illinois, Urbana-Champaign. `http://www.beckman.illinois.edu/events/speechproductionworkshop/index.aspx`.

"Loudness Predicts Prominence; Fundamental Frequency Lends Little", Kochanski, G., Grabe, E., Coleman, J., Rosner, B., 2005 *J. Acoustical Society of America* 11(2), 1038–1054. `http://dx.doi.org/10.1121/1.1923349`.

"Quantitative Measurement of Prosodic Strength in Mandarin", Kochanski, G., Shih, C., Jing, H. 2003 *Speech Communication,* 41(4), 625–645. `http://dx.doi.org/10.1016/S0167-6393(03)00100-6`

"A Quasi-Glottogram signal", Kochanski, G. P. and Shih, Chilin, 2003 *J. Acoustic. Soc. Amer.* 114(4) 2206–2216. `http://dx.doi.org/10.1121/1.1608964`

"Large current density from carbon nanotube field emitters", W. Zhu, C. Bower, O. Zhou, G. Kochanski, S. Jin, *Applied Physics Letters* **75(6)**, August 1999, p. 873–875.

"Faint flickering galaxies: few and far between;" G. P. Kochanski, J. A. Tyson, P. Fischer, 1996 *Astronomical Journal* 111, 1444–1455.

Selected Patents

◇ More than 30 U. S. and/or E. U. Patents granted.

Patent WO/2010/055283: *Acoustic noise prediction and subtraction from intercom audio signal during magnetic resonance imaging*, G. P. Kochanski, D. B. Dobson. Noise cancellation technique that makes use of the signal in the MRI gradient coils to predict and subtract acoustic noise.

U.S. Patent 7,046,812 (2006) *Acoustic beam forming with robust signal estimation* G. P. Kochanski, Man M. Sondhi. Using medians to form beams from an array of microphones for improved rejection of echoes and other out-of-focus sources.

U.S. Patent 8,900,762 *Methods and Apparatus For Location Determination Based On Dispersed Radio Frequency Tags* M. R. Andrews, T. K. Ho, G. P. Kochanski, L. J. Lanzerotti, and D. J. Thomson. How to find your location electronically in a building using cheap passive RFID tags.

U.S. Patent 6,625,576: *Method and apparatus for performing text-to-speech conversion in a client/server environment;* G. P. Kochanski, J. P. Olive, C. Shih. How to transport speech efficiently, both in terms of bandwidth and resources in a cell phone or other portable client.

U.S. Patent 6,310,952: *Telephone Access to Overly Popular Services;* Baldwin, M., Jin, S.-H., Kochanski, G. P. Techniques for auctioning access to telephone services, for instance public golf course reservations.

U.S. Patent 6,219,438: *Produce Identifier using Barcode Scanner and Wavelet Image Processing and having Compensation for Dirt Accumulated on the Viewing Window;* Giordano, D. A., Kochanski, G. P. How can one automatically tell an apple from an orange? Look at the texture.

U.S. Patent 5,483,235: *Stylus-Based Keyboard Arrangement;* G. Kochanski, K. J. Hanson. The sensible way to arrange a keyboard — put common pairs and triplets of letters together, with adjustments for syllable boundaries. We used simulated annealing techniques to optimize the keyboard layout for common messages.

U.S. Patent 7,343,397 (2008): *Method and apparatus for performing predictive cacheing of DNS requests by correlating IP addresses*, G. P. Kochanski. Speeding web access by predicting and pre-fetching DNS records that are likely to be useful.

U.S. Patent 5,512,934: *System and Method for Transmission of Programming on Demand;* G. Kochanski. Describes techniques for distributing video on demand more efficiently. Hardware and spectrum is shared between users by running some users slightly faster, and some slightly slower until their data streams merge.