

Oriol Vinyals
Soda Hall
Berkeley, CA 94709

vinyals@eecs.berkeley.edu
<http://www.cs.berkeley.edu/~vinyals>
510-229-9104

Interests

During my years as a researcher I have developed a deep passion for Artificial Intelligence and Machine Learning, focusing on diverse areas such as Speech Recognition, Speaker Diarization, Computer Vision and Optimization. I am very passionate to be working on these topics as part of my research at Berkeley, and would like to keep doing research on related areas after I graduate.

Education

- **University of California, Berkeley, CA** 2009-2013
Ph.D. student in the EECS department
- **University of California, San Diego, CA** 2008-2009
M.Sc. in the CSE department
- **University of California, Berkeley, CA** 2007-2008
Extension student in the EECS department
- **Polytechnic University of Catalonia, Barcelona, Spain** 2001-2007
Dual degree in Telecommunication Engineering (honors),
and Mathematics (honors)

Experience

- **Research Intern** May 2011 - September 2011
Microsoft Research Redmond, WA
 - research on doing inference on multiparty dialog systems from multimodal data
 - developed a recurrent neural network and a temporal decision tree as inference mechanism
 - Mentors: Dan Bohus and Eric Horvitz
- **Research Intern** May 2010 - August 2010
Microsoft Research Redmond, WA
 - research on the usage of a non linear global feature transformation for speech recognition
 - developed a new optimization scheme faster than traditional stochastic gradient descent
 - Mentor: Daniel Povey
- **Research Intern** June 2009 - August 2009
Microsoft Research Redmond, WA
 - research on the usage of semi-Markov conditional random fields for URL prediction

– Mentor: Patrick Nguyen

- **Research Intern** May 2008 - August 2008
Microsoft Research Redmond, WA
 - research on a better pronunciation modeling via discriminative learning
 - Mentor: Li Deng
- **Research Associate** April 2007 - May 2008
International Computer Science Institute Berkeley, CA
 - research on a better and faster speaker diarization system (who spoke when?)
 - research on a new system that recognizes speakers on the fly
 - development of a faster logarithm function, improving the speed of the standard one
- **Research Assistant (staff)** December 2006 - April 2007
Polytechnic University of Catalonia Barcelona, Spain
 - developed a Catalan voice synthesizer using Festival and HTK toolboxes
- **Research Assistant (staff)** February 2006 - November 2006
The Robotics Institute, CMU Pittsburgh, PA
 - worked on my Final Thesis on the topics of Machine Learning and Computer Vision
 - participated in the CALO project developing a meeting segmentation algorithm based on video information

Academic honors and Scholarships

- Recipient of the 2011 Microsoft Research PhD Fellowship
- Finalist in the Multimedia Grand Challenge of the ACM Multimedia 2010 conference in the Google diaries task, Florence, Italy, 2010
- Finalist of the Qualcomm Innovation Fellowship, 2010
- Finalist of the Best Student Paper Award 2009 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Taipei, Taiwan, 2009
- Awarded with the scholarship of the *Fundacion Caja Madrid*, which covers the costs and living expenses of a 2 year period in any US University, 2007
- Awarded with the European AMIDA training program to do an internship as a visitor researcher in the International Computer Science Institute of the University of California, Berkeley, 2007
- 26 scholarships from the Polytechnic University of Catalonia for obtaining the highest grade in topics such as signal processing, mathematics, computer architecture and networks, 2001 - 2007
- Winner of the Polytechnic University of Catalonia narration challenge, 2003

Publications

- *O. Vinyals, S. Ravuri, D. Povey.*
Revisiting Recurrent Neural Networks for Robust ASR
Proceedings of *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 12)*, Kyoto, Japan, March 2012.
- *O. Vinyals, D. Povey.*
Krylov Subspace Descent for Deep Learning.
In *International Conference on Artificial Intelligence and Statistics (AISTATS 2012)*, Canary Islands, Spain, April 2012.
- *O. Vinyals, D. Povey.*
Krylov Subspace Descent.
In *Neural Information Processing Systems Optimization Workshop and Hierarchical Learning Workshop (NIPS 2011)*, Granada, Spain, December 2011.
- *O. Vinyals, S. Ravuri.*
Comparing Multilayer Perceptron to Deep Belief Network Tandem Features for Robust ASR.
Proceedings of *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 11)*, Prague, Czech Republic, May 2011.
- *X. Anguera, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, O. Vinyals.*
Speaker Diarization: A Review of Recent Research.
IEEE Transactions on Audio, Speech, and Language Processing (TASLP), 2011.
- *G. Friedland, O. Vinyals, T. Darrell.*
Multimodal Location Estimation.
Proceedings of *ACM International Conference on Multimedia (ACM Multimedia 2010)* (full paper), Florence, Italy, October 2010.
- *E. Martin, O. Vinyals, G. Friedland.*
Precise Indoor Localization Using Smart Phones.
Proceedings of *ACM International Conference on Multimedia (ACM Multimedia 2010)* (short paper), Florence, Italy, October 2010.
- *E. Martin, O. Vinyals, G. Friedland, R. Bajcsy.*
Using Android and Indoor Localization for Diaries.
Proceedings of *ACM International Conference on Multimedia (ACM Multimedia 2010)* (ACM MM Grand Challenge), Florence, Italy, October 2010. Finalist of the ACM Multimedia Grand Challenge, Google diaries task.
- *O. Vinyals, E. Martin, G. Friedland.*
Multimodal Indoor Localization: An Audio-Wireless-Based Approach.
Proceedings of *IEEE International Conference on Semantic Computing (ICSC 2010)*, Pittsburgh, PA, September 2010.
- *C. Vaquero, O. Vinyals, G. Friedland.*
A Hybrid Approach to Online Speaker Diarization.
Proceedings of *Interspeech 2010*, Makuhari, Japan, September 2010.

- *S. Bozonnet, N. Evans, X. Anguera, O. Vinyals, G. Friedland, C. Fredouille.*
System Output Combination for Improved Speaker Diarization.
Proceedings of *Interspeech 2010*, Makuhari, Japan, September 2010.
- *O. Vinyals, G. Friedland, N. Morgan.*
Discriminative Training for Hierarchical Clustering in Speaker Diarization.
Proceedings of *Interspeech 2010*, Makuhari, Japan, September 2010.
- *O. Vinyals, L. Deng, D. Yu, A. Acero.*
Discriminative Pronunciation Learning Using Phonetic Decoder and
Minimum-Classification-Error Criterion.
Proceedings of *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 09)*, Taipei, Taiwan, April 2009. Finalist of the Best Student Paper Award.
- *G. Friedland, O. Vinyals, Y. Huang, C. Mller.*
Fusing Short Term and Long Term Features for Improved Speaker Diarization.
Proceedings of *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 09)*, Taipei, Taiwan, April 2009.
- *G. Friedland, O. Vinyals, Y. Huang, C. Mller.*
Prosodic and other Long-Term Features for Speaker Diarization.
In *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, 2009.
- *O. Vinyals, G. Friedland.*
A hardware-independent fast logarithm approximation with adjustable accuracy.
Proceedings of *IEEE International Symposium on Multimedia (ISM 08)*, Berkeley, CA, USA, December 2008.
- *G. Friedland, O. Vinyals.*
Live Speaker Identification in Conversations.
Proceedings of *ACM International Conference on Multimedia (ACM Multimedia 2008)*, Vancouver, Canada, October 2008.
- *O. Vinyals, G. Friedland.*
Modulation Spectrogram Features for Speaker Diarization.
Proceedings of *Interspeech 2008*, Brisbane, Australia, September 2008.
- *K. Boakye, O. Vinyals, G. Friedland.*
Two's a Crowd: Improving Speaker Diarization by Automatically Identifying and Excluding Overlapped Speech.
Proceedings of *Interspeech 2008*, Brisbane, Australia, September 2008.
- *O. Vinyals, G. Friedland.*
Towards semantic analysis of conversations: A system for the live identification of speakers in meetings.
Proceedings of *IEEE International Conference on Semantic Computing (ICSC 08)*, Santa Clara, CA, August 2008.
- *K. Boakye, B. Trueba-Hornero, O. Vinyals, G. Friedland.*
Overlapped speech detection for improved Speaker Diarization in multiparty meetings.
Proceedings of *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 08)*, Las Vegas, NV, April 2008.

- *Y. Huang, O. Vinyals, G. Friedland, C. Mueller, N. Mirghafori, C. Wooters.*
A Fast-Match approach for robust, faster than real-time Speaker Diarization.
Proceedings of *IEEE workshop on Automatic Speech Recognition and Understanding (ASRU 07)*, Kyoto, Japan, December 2007.
- *O. Vinyals, G. Friedland, N. Mirghafori.*
Revisiting a basic function on current CPUs: A fast logarithm implementation with adjustable accuracy.
ICSI TR-07-002, June 2007.
- *F. de la Torre, O. Vinyals.*
Learning Kernel Expansions for Image Classification.
Proceedings of *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 07)*, Minneapolis, MN, June 2007.
- *O. Vinyals, J.M. Gonzalez, S. Blanch, J. Romeu.*
Small antenna design using Shannon's formula.
Proceedings of *IEEE Antennas and Propagation Society International Symposium and USNC/URSI National Radio Science Meeting*. July 2005.

Skills

Languages: Fluent in English. Native speaker of Spanish and Catalan

Programming languages: C/C++, Python, Perl, Java, Assembly

Operating Systems: Linux, Mac OS X, Windows

Applications: MATLAB, Maple, MiniTab, Intel OpenCV, AutoCAD, 3D Studio Max, \LaTeX

Relevant Coursework

@UC Berkeley

- CS281A Statistical Learning Theory
- CS281B Advanced Topics in Learning & Decision Making
- Stat260 Bayesian Modeling and Inference
- CS294 Practical Machine Learning
- CS280 Computer Vision

@UC San Diego

- CSE256 Statistical Natural Language Processing
- ECE273 Convex Optimization and Applications

@CMU

- 10-701 and 15-781 Machine Learning

Professional and personal references furnished upon request.