

Andrei Barbu

(315) 333-0621

andrei@0xab.com
abarbu@csail.mit.edu

January 15, 2019

<http://0xab.com>

Skills & interests

integration of language with robotics and vision, 3D vision, cognitive science, computer vision, functional programming, logic and constraint programming, stochastic and nondeterministic programming languages, linguistics, machine learning, operating systems, neuroimaging to understand the representation of knowledge and language in the brain, video event recognition

Languages

Haskell, Scheme, C, C++, Lisp, Reason, Elm, MATLAB, Java, Python, Clojure, JavaScript
native English and Romanian, basic German and French

Education

Postdoctoral Associate

Massachusetts Institute of Technology, CSAIL

2013 – 2016

Part of the MIT/Harvard Center for Brains, Minds, and Machines
Toyota Research Institute grant *Using vision and language to read minds*
IBM-MIT *Laboratory for Brain-inspired Multimedia Machine Comprehension*
Working on cross-modal reasoning and using vision to solve linguistic tasks.



PhD, Artificial Intelligence

Purdue University, School of Electrical and Computer Engineering

2008 – 2013

Title: Reasoning across language and vision in machines and humans
Advisor: Jeffrey Mark Siskind



language and high-level reasoning for solving inverse vision and manipulation problems
grounding language in vision and robotic manipulation
stochastic modeling via probabilistic programming
optimization of large stochastic cognitive models
segmentation and tracking of objects in videos with high-level semantic priors
large-scale video search using sentential queries
the neural representation of verbs, events, and compositionality
<http://0xab.com/research/>

Bachelor of Computer Science

University of Waterloo

2004 – 2008

complexity theory, real-time operating systems, graph theory, algorithm design,
programming languages, medical imaging, linear programming, coding theory



Employment History

Research Scientist

Massachusetts Institute of Technology

Oct. 2016 – present

Cambridge, USA

Part of the MIT/Stanford Toyota Research Institute (TRI)
Part of the MIT/Harvard Center for Brains, Minds, and Machines (CBMM)
Part of the MIT/IBM Laboratory for Brain-inspired Multimedia Machine
Comprehension (BM³C)
Developing robots that learn, understand, and follow commands
Working on deep video understanding grounding language in perception
Teaching and co-organizing the CBMM summer school (35 graduate students, 40 speakers)



Postdoctoral Associate

Advisor: Boris Katz

Part of the MIT/Harvard Center for Brains, Minds, and Machines (CBMM)
Disambiguating language in the context of vision
Language acquisition and grounded question answering
Teaching and co-organizing the CBMM summer school

Dec. 2013 – Oct. 2016

MIT, USA

**Research Assistant**

Advisor: Jeffrey Mark Siskind

solved computer-vision and manipulation problems using AD-based optimization
designed and built custom robots
implemented stochastic programs to use and produce language
participated in the DARPA *Mind's Eye* program which investigated event recognition
won the yearly evaluations against 11 other teams in both years of the program
developed new robust segmentation and tracking methods for people and objects
processed millions of frames of video on a top-500 supercomputer

Jan. 2008 – Nov. 2013

Purdue University, USA

**Software Development Engineer**

Amazon.com

migrated from a monolithic database design to a distributed service-oriented system
designed & developed distributed services to collect metrics about internal systems
coordinated with other teams to develop internal tools and alarms

May 2010 – Aug. 2010

Seattle, USA

**Kernel Developer**

Google Summer of Code 2008 – the GNU Hurd

designed and implemented an instrumentation framework for GNU Mach
implemented the DWARF2 debugging standard
debugger improvements
bug fixes and other new features

May 2008 – Aug. 2008

Purdue University, USA

Google
Summer of Code**Undergraduate Research Assistant**

Advisor: Professor Jeffrey Mark Siskind

numerous improvements to a C frontend, including updates for C99 and GNUC
improvements to the existing Scheme compiler
ported the Scheme→C compiler to x86-64
ported various other applications to x86-64
significant performance improvements

July 2006 – Sept. 2006

Purdue University, USA

**Software Developer and Knowledge Systems Architect**

Dalin Software SA

ontology, temporal, and spatial logic frameworks
designed extensions to OWL, the Ontology Web Language
designed and developed tools for querying and updating OWL+SWRL ontologies

July 2005 – May 2006

Laussane, Switzerland

Publications

- Safeguarding the visual world for those with photosensitive epilepsy *in review.*
D. Banda, A. Barbu, and B. Katz P1
Pattern Recognition Letters (PatRec)
- Partially Occluded Hands: A new dataset for single-image hand pose estimation *December 2018*
B. Myanganbayar, C. Mata, G. Dekel, B. Katz, G. Ben-Yosef, A. Barbu P2
Asian Conference on Computer Vision (ACCV)
- Grounding semantic parsing using caption videos *November 2018*
C. Ross, A. Barbu, Yevgeni Berzak, Battushig Myanganbayar, and B. Katz P3
Empirical Methods in Natural Language Processing (EMNLP)
- Deep sequential models for sampling-based planning *October 2018*
Y. Kuo, A. Barbu, and B. Katz P4
International Conference on Intelligent Robots (IROS)
<http://0xab.com/papers/iros2018.pdf>
- Deep compositional models for robotic planning and language *October 2018*
Y. Kuo, A. Barbu, and B. Katz P5
Workshop on Language and Robotics at IROS
- Temporal Grounding Graphs for Language Understanding with Accrued Visual-Linguistic Context *August 2017*
R. Paul, A. Barbu, S. Felshin, B. Katz, and N. Roy P6
International Joint Conference on Artificial Intelligence (IJCAI)
<http://0xab.com/papers/ijcai17.pdf>
- Temporal Grounding Graphs for Language Understanding with Accrued Visual-Linguistic Context *July 2017*
R. Paul, A. Barbu, S. Felshin, B. Katz and N. Roy P7
Language Grounding for Robotics Workshop at ACL
<http://0xab.com/papers/ac17.pdf>
- Saying What You're Looking For: Linguistics Meets Video Search *October 2016*
D. Barrett, A. Barbu, N. Siddharth, and J. M. Siskind P8
IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)
<http://0xab.com/papers/pami.pdf>
- Anchoring and Agreement in Syntactic Annotations *September 2016*
Y. Berzak, Y. Huang, A. Barbu, A. Korhonen, and B. Katz P9
Empirical Methods in Natural Language Processing (EMNLP)
<https://arxiv.org/pdf/1605.04481v3>
- Do You See What I Mean? Visual Resolution of Linguistic Ambiguities *September 2015*
Y. Berzak, A. Barbu, D. Harari, B. Katz, and S. Ullman P10
Empirical Methods in Natural Language Processing (EMNLP)
<http://0xab.com/papers/emnlp2015.pdf>
- A Compositional Framework for Grounding Language Inference, Generation, and Acquisition in Video *April 2015*
H. Yu, N. Siddharth, A. Barbu, and J. M. Siskind P11
Journal of Artificial Intelligence Research (JAIR)
<http://0xab.com/papers/jair.pdf>
- Seeing is Worse than Believing: Reading People's Minds Better than Computer-Vision Methods Recognize Actions *September 2014*
A. Barbu, D. P. Barrett, W. Chen, N. Siddharth, C. Xiong, J. J. Corso, C. D. Fellbaum, C. Hanson, S. J. Hanson, S. Hélie, E. Malaia, B. A. Pearlmutter, J. M. Siskind, T. M. Talavage, and R. B. Wilbur P12
European Conference on Computer Vision (ECCV)
<http://0xab.com/papers/eccv2014.pdf>

- Seeing What You're Told: Sentence-Guided Activity Recognition In Video** September 2014
 N. Siddharth, A. Barbu, J. M. Siskind P13
Human-Machine Communication for Visual Recognition and Search Workshop at ECCV
<ftp://ftp.ecn.purdue.edu/qobi/cvpr2014.pdf>
- Language-driven video retrieval** June 2014
 A. Barbu, N. Siddharth, J. M. Siskind P14
Vision Meets Cognition at CVPR
http://www.visionmeetscognition.org/fpic2014/Camera_Ready/Paper%2033.pdf
- Seeing what you're told: sentence-guided activity recognition in video** June 2014
 N. Siddharth, A. Barbu, and J. M. Siskind P15
IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
<http://0xab.com/papers/cvpr2014.pdf>
- Recognizing human activities from partially observed videos** June 2013
 Y. Cao, D. Barrett, A. Barbu, N. Siddharth, H. Yu, A. Michaux, Y. Lin, S. Dickinson, J. M. Siskind, and S. Wang P16
IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
 poster 472/1870 (25.2%)
<http://0xab.com/papers/cvpr2013.pdf>
- Seeing unseeability to see the unseeable** December 2012
 N. Siddharth, A. Barbu, and J. M. Siskind P17
Advances in Cognitive Systems, 2:77–94
 oral 14/38 (37%), conference presentation associated with
<http://0xab.com/papers/acs2012a.pdf>
- Simultaneous object detection, tracking, and event recognition** December 2012
 A. Barbu, N. Siddharth, A. Michaux, and J. M. Siskind P18
Advances in Cognitive Systems, 2:203–20
 oral 14/38 (37%), conference presentation associated with
<http://0xab.com/papers/acs2012b.pdf>
- Video in sentences out** August 2012
 A. Barbu, A. Bridge, Z. Burchill, D. Coroian, S. Dickinson, S. Fidler, A. Michaux, S. Mussman, N. Siddharth, D. Salvi, L. Schmidt, J. Shangguan, J. M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang P19
Conference on Uncertainty In Artificial Intelligence (UAI), pp. 102–12
 oral 24/304 (8%)
<http://0xab.com/papers/uai2012.pdf>
- A visual language model for estimating object pose and structure in a generative visual domain** May 2011
 N. Siddharth, A. Barbu, and J. M. Siskind P20
IEEE International Conference on Robotics and Automation (ICRA)
 oral 982/2004 (49%)
<http://0xab.com/papers/icra2011.pdf>
- Learning physically-instantiated robotic game play through visual observation** May 2010
 A. Barbu, N. Siddharth, and J. M. Siskind P21
IEEE International Conference on Robotics and Automation (ICRA)
 oral 856/2062 (42%)
<http://0xab.com/papers/icra2010.pdf>

Patents

- Safeguarding the perceptual world for those with photosensitivity** *in prep.*
A. Barbu, D. Banda, B. Katz A1
Preliminary filed, patent number to be assigned
- Correlating videos and sentences** *November 2015*
J. M. Siskind, A. Barbu, N. Siddharth, and H. Yu A2
US patent US9183466

Technical Reports

- Bias and Agreement in Syntactic Annotations** *May 2016*
Yevgeni Berzak, Yan Huang, A. Barbu, Anna Korhonen, Boris Katz T1
arXiv:1605.04481
<http://arxiv.org/abs/1605.04481>
- The Compositional Nature of Event Representations in the Human Brain** *June 2013*
A. Barbu, N. Siddharth, C. Xiong, J. J. Corso, C. D. Fellbaum, C. Hanson, S. J. Hanson, S. Hélie, E. Malaia, B. A. Pearlmutter, J. M. Siskind, T. M. Talavage, and R. B. Wilbur T2
arXiv:1306.2293
<http://arxiv.org/abs/1306.2293>
- Large-scale automatic labeling of video events with verbs based on event-participant interaction** *April 2012*
A. Barbu, A. Bridge, D. Coroian, S. Dickinson, S. Mussman, S. Narayanaswamy, D. Salvi, L. Schmidt, J. Shangguan, J. M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang T3
arXiv:1204.3616
<http://arxiv.org/abs/1204.3616>

Workshops organized

- Combining Vision and Language ACCV 2018** *June 2018*
Qi Wu, A. Barbu W1
to appear
accv2018.net/program
- Vision and Language at CVPR 2018** *June 2018*
A. Barbu, Siddharth Narayanaswamy, Tao Mei, Yevgeni Berzak, Nishant Shukla, Jiebo Luo, Rahul Sukthankar W2
16 accepted papers, 7 invited speakers, over 200 attendees
languageandvision.com
- Vision and Language at CVPR 2017** *July 2017*
A. Barbu, Tao Mei, Siddharth Narayanaswamy, Puneet Kumar Dokania, Quanshi Zhang, Nishant Shukla, Jiebo Luo, Rahul Sukthankar W3
17 accepted papers, 7 invited speakers, over 200 attendees
languageandvision.com/2017.html
- Vision and Language at CVPR 2015** *June 2015*
A. Barbu, Georgios Evangelopoulos, Daniel Harari, Krystian Mikolajczyk, Siddharth Narayanaswamy, Caiming Xiong, Yibiao Zhang W4
14 accepted papers, 9 invited speakers, over 100 attendees
languageandvision.com/2015.html

Teaching

CBMM Summer school 2018

Gabriel Kreiman, Tomaso Poggio, Boris Katz

One of two main organizers for a three week summer school for 35 undergraduates, graduate students, and postdocs with 40 speakers and 15 teaching assistants. Developed the schedule, selected speakers and topics, lectured on several subjects, and supervised several projects.

cbmm.mit.edu/summer-school

August 2018

E1

CBMM Summer school 2017

Tomaso Poggio, Gabriel Kreiman

One of two main organizers for a three week summer school for 35 undergraduates, graduate students, and postdocs with 40 speakers and 15 teaching assistants. Developed the schedule, selected speakers and topics, lectured on several subjects, and supervised several projects.

cbmm.mit.edu/summer-school

August 2017

E2

CBMM Summer school 2016

Tomaso Poggio, Gabriel Kreiman

One of two main organizers for a three week summer school for 35 undergraduates, graduate students, and postdocs with 40 speakers and 15 teaching assistants. Developed the schedule, selected speakers and topics, lectured on several subjects, and supervised several projects.

cbmm.mit.edu/summer-school

August 2016

E3

CBMM Summer school 2015

Tomaso Poggio, Gabriel Kreiman, L. Mahadevan

One of two main organizers for a three week summer school for 30 undergraduates, graduate students, and postdocs with over 30 speakers and 10 teaching assistants. Developed the schedule, selected speakers and topics, lectured on several several subjects, and supervised several projects.

cbmm.mit.edu/summer-school

August 2015

E4

CBMM Summer school 2014

Tomaso Poggio, Gabriel Kreiman, L. Mahadevan

Helped organize and teach a two week summer school with 20 undergraduate and graduate students. Lectured, helped develop the schedule, and managed projects.

cbmm.mit.edu/summer-school

June 2014

E5

Teaching assistant for ECE 570 Artificial Intelligence (unofficial)

Jeffrey Mark Siskind

TA for four offerings of the graduate AI course. TA duties were part of the RA and not formally listed; they included office hours and grading assignments.

<https://engineering.purdue.edu/~ee570>

2008-2012

E6

Teaching assistant for ECE 473 Introduction to Artificial Intelligence (unofficial)

Jeffrey Mark Siskind

TA for four offerings of the undergraduate introduction to AI course. TA duties were part of the RA and not formally listed; they included office hours and grading assignments.

<https://engineering.purdue.edu/~ee473>

2008-2012

E7

PhD students supervised

Adam Yaari

Language and neuroscience

Co-supervised with Boris Katz.

2018-present

S1

Yen-Ling Kuo

Grounding language in planning

Co-supervised with Boris Katz.

2017-present

S2

Jonathan Malmaud

Grounding language in physics

Co-supervised with Boris Katz and Josh Tenenbaum.

2017-present

S3

Candace Ross 2016-present
 Grounded language acquisition
Co-supervised with Boris Katz. S4

Master's students supervised

Julian Alverio *current*
 Robots that understand language
Co-supervised with Boris Katz. M1

David Mayo *current*
 Understanding human and machine object recognition at scale
Co-supervised with Boris Katz. M2

Dalitso Banda 2018
 Removing dangerous visual stimuli from videos
Co-supervised with Boris Katz. M3

Battushig Myanganbayar 2018
 Inverse graphics for hand-pose reconstruction
Co-supervised with Boris Katz now at Apple. M4

Maria Ryskina 2016
 Imaginative Paraphrase Recognition via Join Vision-Language Video Analysis
Co-supervised with Boris Katz. Now a PhD student at CMU. M5

Sergey Voronov 2016
 Common Sense Reasoning Through Imagination
Co-supervised with Boris Katz. Now a PhD student at University of Ohio. M6

Nicolas Rakover 2016
 A uniform Representation for Visual Concepts
Co-supervised with Boris Katz. Now at Google. M7

Undergraduate students supervised

Nazar Ilamanov *present*
Co-supervised with Boris Katz. U1

Julian Alverio 2018
Co-supervised with Boris Katz. Now a Masters student at MIT. U2

Gil Dekel 2016-2018
Co-supervised with Boris Katz. Now at Google. U3

Battushig Myanganbayar 2016
Co-supervised with Boris Katz. Now a Masters student at MIT. U4

David Mayo 2016
Co-supervised with Boris Katz. Now a Masters student at MIT. U5

Corinn Herrik 2016
Co-supervised with Boris Katz. U6

Sarah Coe 2015-2016
Co-supervised with Boris Katz. U7

Candace Ross 2015
Co-supervised with Boris Katz. Now a PhD student at MIT. U8

Victor Cabrera 2015
Co-supervised with Boris Katz. U9

Matthew Evanusa 2014
Co-supervised with Boris Katz. U10

Matthew Arbesfeld 2014-2015
Co-supervised with Boris Katz. U11