

Andrei Barbu

(315) 333-0621

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

July 25, 2017
<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++, Clojure, Lisp, C, Prolog, MATLAB, Java, Python, Julia, JavaScript, SQL,
x86 and MIPS assembly
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*
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



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

- Temporal Grounding Graphs for Language Understanding with Accrued Visual-Linguistic Context
R. Paul, A. Barbu, S. Felshin, B. Katz and N. Roy
International Joint Conference on Artificial Intelligence (IJCAI)
<http://0xab.com/papers/ijcai17.pdf> August 2017 1
- Planning with Language Grounded in Perception
C. Ross, A. Barbu, B. Katz
WiNLP (Women in NLP) Workshop at ACL July 2017 2
- Temporal Grounding Graphs for Language Understanding with Accrued Visual-Linguistic Context
R. Paul, A. Barbu, S. Felshin, B. Katz and N. Roy
Language Grounding for Robotics Workshop at ACL
<http://0xab.com/papers/acl17.pdf> July 2017 3
- Saying What You're Looking For: Linguistics Meets Video Search
D. Barrett, A. Barbu, N. Siddharth, and J. M. Siskind
IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)
<http://0xab.com/papers/pami.pdf> October 2016 4
- Anchoring and Agreement in Syntactic Annotations
Yevgeni Berzak, Yan Huang, A. Barbu, Anna Korhonen, Boris Katz
Empirical Methods in Natural Language Processing (EMNLP)
<https://arxiv.org/pdf/1605.04481v3> September 2016 5
- Do You See What I Mean? Visual Resolution of Linguistic Ambiguities
Yevgeni Berzak, A. Barbu, Daniel Harari, Boris Katz, and Shimon Ullman
Empirical Methods in Natural Language Processing (EMNLP)
<http://0xab.com/papers/emnlp2015.pdf> September 2015 6
- A Compositional Framework for Grounding Language Inference, Generation, and Acquisition in Video
H. Yu, N. Siddharth, A. Barbu, and J. M. Siskind
Journal of Artificial Intelligence Research (JAIR)
<http://0xab.com/papers/jair.pdf> April 2015 7
- Seeing is Worse than Believing: Reading People's Minds Better than Computer-Vision Methods Recognize Actions
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
European Conference on Computer Vision (ECCV)
<http://0xab.com/papers/eccv2014.pdf> September 2014 8
- Seeing What You're Told: Sentence-Guided Activity Recognition In Video
N. Siddharth, A. Barbu, J. M. Siskind
Human-Machine Communication for Visual Recognition and Search Workshop at ECCV
<ftp://ftp.ecn.purdue.edu/qobi/cvpr2014.pdf> September 2014 9
- Language-driven video retrieval
A. Barbu, N. Siddharth, J. M. Siskind
Vision Meets Cognition at CVPR
http://www.visionmeetscognition.org/fpic2014/Camera_Ready/Paper%2033.pdf June 2014 10
- Seeing what you're told: sentence-guided activity recognition in video
N. Siddharth, A. Barbu, and J. M. Siskind
IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
<http://0xab.com/papers/cvpr2014.pdf> June 2014 11

	June 2013
<p>Recognizing human activities from partially observed videos Y. Cao, D. Barrett, A. Barbu, N. Siddharth, H. Yu, A. Michaux, Y. Lin, S. Dickinson, J. M. Siskind, and S. Wang <i>IEEE Conference on Computer Vision and Pattern Recognition (CVPR)</i> poster 472/1870 (25.2%) http://0xab.com/papers/cvpr2013.pdf</p>	12
	December 2012
<p>Seeing unseeability to see the unseeable N. Siddharth, A. Barbu, and J. M. Siskind <i>Advances in Cognitive Systems</i>, 2:77–94 oral 14/38 (37%), conference presentation associated with http://0xab.com/papers/acs2012a.pdf</p>	13
	December 2012
<p>Simultaneous object detection, tracking, and event recognition A. Barbu, N. Siddharth, A. Michaux, and J. M. Siskind <i>Advances in Cognitive Systems</i>, 2:203–20 oral 14/38 (37%), conference presentation associated with http://0xab.com/papers/acs2012b.pdf</p>	14
	August 2012
<p>Video in sentences out 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 <i>Conference on Uncertainty In Artificial Intelligence (UAI)</i>, pp. 102–12 oral 24/304 (8%) http://0xab.com/papers/uai2012.pdf</p>	15
	May 2011
<p>A visual language model for estimating object pose and structure in a generative visual domain N. Siddharth, A. Barbu, and J. M. Siskind <i>IEEE International Conference on Robotics and Automation (ICRA)</i> oral 982/2004 (49%) http://0xab.com/papers/icra2011.pdf</p>	16
	May 2010
<p>Learning physically-instantiated robotic game play through visual observation A. Barbu, N. Siddharth, and J. M. Siskind <i>IEEE International Conference on Robotics and Automation (ICRA)</i> oral 856/2062 (42%) http://0xab.com/papers/icra2010.pdf</p>	17

Patents

	November 2015
<p>Correlating videos and sentences J. M. Siskind, A. Barbu, N. Siddharth, and H. Yu <i>US patent US9183466</i></p>	18

Technical Reports

	May 2016
<p>Bias and Agreement in Syntactic Annotations Yevgeni Berzak, Yan Huang, A. Barbu, Anna Korhonen, Boris Katz <i>arXiv:1605.04481</i> http://arxiv.org/abs/1605.04481</p>	19
	June 2013
<p>The Compositional Nature of Event Representations in the Human Brain 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 <i>arXiv:1306.2293</i> http://arxiv.org/abs/1306.2293</p>	20

April 2012

Large-scale automatic labeling of video events with verbs
based on event-participant interaction

A. Barbu, A. Bridge, D. Coroian, S. Dickinson, S. Mussman, S. Narayanaswamy, D. Salvi, L. Schmidt, J. Shangquan, J. M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang
arXiv:1204.3616
<http://arxiv.org/abs/1204.3616>

21

Workshops organized

July 2017

Vision and Language at CVPR 2017

A. Barbu, Tao Mei, Siddharth Narayanaswamy, Puneet Kumar Dokania, Quanshi Zhang, Nishant Shukla, Jiebo Luo, Rahul Sukthankar
17 accepted papers, 7 invited speakers, over 200 attendees
languageandvision.com/2017.html

22

June 2015

Vision and Language at CVPR 2015

A. Barbu, Georgios Evangelopoulos, Daniel Harari, Krystian Mikolajczyk, Siddharth Narayanaswamy, Caiming Xiong, Yibiao Zhang
14 accepted papers, 9 invited speakers, over 100 attendees
languageandvision.com/2015.html

23

Teaching

August 2017

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

24

August 2016

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

25

August 2015

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

26

June 2014

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

27

2008-2012

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>

28

2008-2012

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>

29

PhD students supervised

Yen-Ling Kuo Grounding language in planning <i>Co-supervised with Boris Katz.</i>	2016-present 30
Candace Ross Grounded language acquisition <i>Co-supervised with Boris Katz.</i>	2016-present 31

Master's students supervised

David Mayo Understanding human and machine object recognition at scale <i>Co-supervised with Boris Katz.</i>	current 32
Maria Ryskina Imaginative Paraphrase Recognition via Joint Vision-Language Video Analysis <i>Co-supervised with Boris Katz. Now a PhD student at CMU.</i>	2016 33
Sergey Voronov Common Sense Reasoning Through Imagination <i>Co-supervised with Boris Katz. Now a PhD student at University of Oregon.</i>	2016 34
Nicolas Rakover A uniform Representation for Visual Concepts <i>Co-supervised with Boris Katz. Now at Google.</i>	2016 35

Undergraduate students supervised

Nazar Ilamanov <i>Co-supervised with Boris Katz.</i>	present 36
Battushig Myanganbayar <i>Co-supervised with Boris Katz. Now a Masters student at MIT.</i>	2016 37
David Mayo <i>Co-supervised with Boris Katz. Now a Masters student at MIT.</i>	2016 38
Corinn Herrik <i>Co-supervised with Boris Katz.</i>	2016 39
Sarah Coe <i>Co-supervised with Boris Katz.</i>	2015-2016 40
Candace Ross <i>Co-supervised with Boris Katz. Now a PhD student at MIT.</i>	2015 41
Victor Cabrera <i>Co-supervised with Boris Katz.</i>	2015 42
Matthew Evanusa <i>Co-supervised with Boris Katz.</i>	2014 43
Matthew Arbesfeld <i>Co-supervised with Boris Katz.</i>	2014-2015 44