

Andrei Barbu

(315) 333-0621

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

June 29, 2023
<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, machine learning for accessibility and fairness

Languages

Haskell, Scheme, C, C++, Lisp, 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 grant *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
TA for ECE473 and ECE570, Artificial Intelligence
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
processed tens of millions of frames of video on a top-500 supercomputer
solved computer-vision and manipulation problems using AD-based optimization
<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)

Part of the Air Force-MIT AI Accelerator

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)

ONR award *Developing next generation AI vision systems by characterizing and exploiting untapped primate visual processing circuit motifs*

NSF award *NCS-FO: Studying language in the brain in the modern machine learning era*

DARPA award *Grounded Artificial Intelligence Language Acquisition (GAILA)*

DARPA award *Knowledge Management at Scale and Speed (KMASS)*

Software Development Engineer

Amazon.com

May 2010 – Aug. 2010

Seattle, USA



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

Kernel Developer

Google Summer of Code 2008 – the GNU Hurd

May 2008 – Aug. 2008

Purdue University, USA



Google
Summer of Code

designed and implemented an instrumentation framework for GNU Mach

implemented the DWARF2 debugging standard

debugger improvements

bug fixes and other new features

Undergraduate Research Assistant

Advisor: Professor Jeffrey Mark Siskind

July 2006 – Sept. 2006

Purdue University, USA



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

Software Developer and Knowledge Systems Architect

Dalin Software SA

July 2005 – May 2006

Laussane, Switzerland

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

Publications

- [BrainBERT: Self-supervised representation learning for intracranial recordings](#) March 2023
C. Wang, V. Subramaniam, A. Yaari, G. Kreiman, B. Katz, I. Cases, A. Barbu P1
International Conference on Learning Representations (ICLR)
- [Zero-shot linear combinations of grounded social interactions with Linear Social MDPs](#) January 2023
R. Tejwani, Y. Kuo, T. Shu, B. Stankovits, D. Gutfreund, J. Tenenbaum, B. Katz, A. Barbu P2
Conference on Artificial Intelligence (AAAI)
- [The Aligned Multimodal Movie Treebank: An audio, video, dependency-parse treebank](#) September 2022
A. Yaari, J. DeWitt, H. Hu, B. Stankovits, S. Felshin, Y. Berzak, H. Aparicio, B. Katz, I. Cases, A. Barbu P3
Empirical Methods in Natural Language Processing (EMNLP)
- [Developing a Series of AI Challenges for the United States Department of the Air Force](#) September 2022
V. Gadepally, G. Angelides, A. Barbu, A. Bowne, L. Brattain, T. Broderick, A. Cabrera, G. Carl, R. Carter, M. Cha, E. Cowen, J. Cummings, B. Freeman, J. Glass, S. Goldberg, M. Hamilton, T. Heldt, K. W. Huang, P. Isola, B. Katz, J. Koerner, Y. Lin, D. Mayo, K. McAlpin, T. Perron, J. Piou, H. M. Rao, H. Reynolds, K. Samuel, S. Samsi, M. Schmidt, L. Shing, O. Simek, B. Swenson, V. Sze, J. Taylor, P. Tylkin, M. Veillette, M. L. Weiss, A. Wollaber, S. Yuditskaya, J. Kepner P4
IEEE Conference on High Performance Extreme Computing (HPEC)
- [Quantifying the Emergence of Symbolic Communication](#) August 2022
E. Cheng, Y. Kuo, I. Cases, B. Katz, A. Barbu P5
CogSci
- [Incorporating rich social interactions into MDPs](#) March 2022
R. Tejwani, Y. Kuo, T. Shu, B. Stankovits, D. Gutfreund, J. Tenenbaum, B. Katz, A. Barbu P6
International Conference on Robotics and Automation (ICRA)
Excellent Paper award at IROS Cognitive and Social Aspects of Human Multi-Robot Interaction Workshop
- [Trajectory prediction with linguistic representations](#) March 2022
Y. Kuo, X. Huang, A. Barbu, S. McGill, B. Katz, J. Leonard, G. Rosman P7
International Conference on Robotics and Automation (ICRA)
- [Neural Regression, Representational Similarity, Model Zoology & Neural Taskonomy at Scale in Rodent Visual Cortex](#) December 2021
C. Conwell, D. Mayo, A. Barbu, M. Buice, G. Alvarez, B. Katz P8
Conference on Neural Information Processing Systems (NeurIPS)
- [Perceiving social events in a physical world](#) September 2021
T. Shu, A. Netanyahu, M. Kryven, J. Muchovej, N. Shenoy, B. Katz, A. Barbu, T. Ullman, J. Tenenbaum P9
Vision Sciences Society Annual Meeting
- [Social Interactions as Recursive MDPs](#) August 2021
R. Tejwani, Y. Kuo, B. Katz, and A. Barbu P10
Conference on Robot Learning (CoRL)
- [Compositional Networks Enable Systematic Generalization for Grounded Language Understanding](#) August 2021
Y. Kuo, B. Katz, and A. Barbu P11
Findings of Empirical Methods in Natural Language Processing (EMNLP)
- [Measuring Social Biases in Grounded Vision and Language Embeddings](#) June 2021
C. Ross, B. Katz, and A. Barbu P12
Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)

| | |
|--|-------------------------------|
| <p>Multi-resolution modeling of a discrete stochastic process identifies causes of cancer A. Yaari, M. Sherman, O. Clarke Priebe, P. Loh, B. Katz, A. Barbu, B. Berger <i>International Conference on Learning Representations (ICLR)</i></p> | <p>May 2021 P13</p> |
| <p>PHASE: PHYSically-grounded Abstract Social Events for Machine Social Perception A. Netanyahu, T. Shu, B. Katz, A. Barbu, J. Tenenbaum <i>Conference on Artificial Intelligence (AAAI)</i> Best paper award at the NeurIPS SVRHM workshop</p> | <p>January 2021 P14</p> |
| <p>Genetic and epigenetic factors interact to explain somatic mutation rates at multiple length scales A. Yaari, M. Sherman, O. Clarke Priebe, P. Loh, A. Barbu, B. Katz, B. Berger <i>American Society of Human Genetics</i> Best paper award semi-finalist and reviewer's choice award https://www.abstractsonline.com/pp8/#!/9070/presentation/3409</p> | <p>October 2020 P15</p> |
| <p>Learning a natural-language to LTL executable semantic parser for grounded robotics C. Wang, C. Ross, Y. Kuo, B. Katz, and A. Barbu <i>Conference on Robot Learning (CoRL)</i></p> | <p>November 2020 P16</p> |
| <p>Encoding formulas as deep networks: Reinforcement learning for zero-shot execution of LTL formulas Y. Kuo, B. Katz, and A. Barbu <i>International Conference on Intelligent Robots (IROS)</i></p> | <p>October 2020 P17</p> |
| <p>Deep compositional robotic planners that follow natural language commands Y. Kuo, B. Katz, and A. Barbu <i>International Conference on Robotics and Automation (ICRA)</i></p> | <p>May 2020 P18</p> |
| <p>ObjectNet: A large-scale bias-controlled dataset for pushing the limits of object recognition models A. Barbu, D. Mayo, J. Alverio, W. Luo, C. Wang, D. Gutfreund, J. Tenenbaum, and B. Katz <i>Neural Information Processing Systems (NeurIPS)</i> https://objectnet.dev</p> | <p>December 2019 P19</p> |
| <p>Deep video-to-video transformations for accessibility with an application to photosensitivity D. Banda, A. Barbu, and B. Katz <i>Pattern Recognition Letters (PatRec)</i></p> | <p>June 2019 P20</p> |
| <p>Partially Occluded Hands: A new dataset for single-image hand pose estimation B. Myanganbayar, C. Mata, G. Dekel, B. Katz, G. Ben-Yosef, A. Barbu <i>Asian Conference on Computer Vision (ACCV)</i></p> | <p>December 2018 P21</p> |
| <p>First steps to understanding how symbolic communication arises J. Correa, Y. Kuo, B. Katz, and A. Barbu <i>Workshop on Emergent Communication at NeurIPS</i></p> | <p>December 2018 P22</p> |
| <p>Grounding semantic parsing using caption videos C. Ross, A. Barbu, Yevgeni Berzak, Battushig Myanganbayar, and B. Katz <i>Empirical Methods in Natural Language Processing (EMNLP)</i> http://aclweb.org/anthology/D18-1285</p> | <p>November 2018 P23</p> |
| <p>Deep sequential models for sampling-based planning Y. Kuo, A. Barbu, and B. Katz <i>International Conference on Intelligent Robots (IROS)</i> http://0xab.com/papers/iros2018.pdf</p> | <p>October 2018 P24</p> |
| <p>Deep compositional models for robotic planning and language Y. Kuo, A. Barbu, and B. Katz <i>Workshop on Language and Robotics at IROS</i></p> | <p>October 2018 P25</p> |

Temporal Grounding Graphs for Language Understanding with Accrued Visual-Linguistic Context August 2017
R. Paul, A. Barbu, S. Felshin, B. Katz, and N. Roy P26
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 P27
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 P28
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 P29
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 P30
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 P31
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 P32
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 P33
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 P34
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 P35
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 P36
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 P37
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 P38
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 P39
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 P40
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 P41
IEEE International Conference on Robotics and Automation (ICRA)
 oral 856/2062 (42%)
<http://0xab.com/papers/icra2010.pdf>

Patents

- Deep Compositional Robotic Planners That Follow Natural Language Commands** in review
 Y. Kuo, Boris Katz, A. Barbu A1
Provisional 62/944,924 and 62/944,932
- Computer method and apparatus making screens safe for those with photosensitivity** July 2022
 A. Barbu, D. Banda, B. Katz A2
US patent US11381715
- Correlating videos and sentences** November 2015
 J. M. Siskind, A. Barbu, N. Siddharth, and H. Yu A3
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

Linguistics Meets Image and Video Retrieval at ICCV 2019 October 2019
Amrita Saha, Hui Wu, Adriana I. Kovashka, A. Barbu, Xiaoxiao Guo, Karthik Sankaranarayanan, Samarth Bharadwaj, and Yupeng Gao W1
3 invited speakers, hosted the Fashion IQ challenge
<https://sites.google.com/view/lingir>

Vision and Language at CVPR 2019 June 2019
A. Barbu, Siddharth N., D. Gutfreund W2
22 accepted papers, 5 invited speakers, over 200 attendees
languageandvision.com

Combining Vision and Language ACCV 2018 June 2018
Qi Wu, Peng Wang, Chuang Gan, Fumin Shen, Xiaodong He, A. Barbu, Anton van den Hengel W3
A workshop introducing new researchers to problems in vision and language, 7 invited speakers
http://qi-wu.me/accv_v21/

Vision and Language at CVPR 2018 June 2018
A. Barbu, Siddharth N., Tao Mei, Y. Berzak, N. Shukla, J. Luo, R. Sukthankar, D. Gutfreund W4
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 W5
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 W6
14 accepted papers, 9 invited speakers, over 100 attendees
languageandvision.com/2015.html

PhD students supervised

Christopher Arnold 2023-present
Multimodal story understanding S1
Co-supervised with Boris Katz.

Ravi Tejwani 2022-2023
Social perception S2
Co-supervised with Boris Katz.

Christopher Z Wang 2021-present
Object recognition S3
Co-supervised with Boris Katz.

David Mayo 2020-present
Object recognition S4
Co-supervised with Boris Katz.

| | |
|---|--------------------|
| Aviv Netanyahu Social interactions <i>Co-supervised with Boris Katz.</i> | 2019-2020 S5 |
| Adam Yaari Language and neuroscience <i>CEO and founder of Serinus Biosciences. Co-supervised with Boris Katz.</i> | 2018-present S6 |
| Yen-Ling Kuo Grounding language in planning <i>Faculty at University of Virginia. Co-supervised with Boris Katz.</i> | 2017-2022 S7 |
| Candace Ross Grounded language acquisition <i>Now a Postdoc at FAIR. Co-supervised with Boris Katz.</i> | 2016-2022 S8 |

Master's students supervised

| | |
|---|---------------|
| Abdulrahman Alabdulkareem Understanding modalities in detail <i>Co-supervised with Boris Katz.</i> | current M1 |
| Anastasiia Uvarova Modeling the development of language <i>Co-supervised with Boris Katz.</i> | current M2 |
| Laura Queipo Causal summarization <i>Co-supervised with Boris Katz.</i> | current M3 |
| Vighnesh Subramaniam Understanding multimodal integration in the brain <i>Co-supervised with Boris Katz.</i> | current M4 |
| Meagan R. Jens Machines that change the rules to win <i>Co-supervised with Boris Katz.</i> | 2023 M5 |
| Dana Rosenfarb Creating a large-scale benchmark for language in the brain <i>Co-supervised with Boris Katz.</i> | 2022 M6 |
| Felipe Monsalve Developing a secure ML platform for healthcare applications <i>Co-supervised with Boris Katz.</i> | 2022 M7 |
| Emily Cheng Understanding language evolution <i>Now a PhD student. Co-supervised with Boris Katz.</i> | 2022 M8 |
| Aaditya K. Singh Deep Attentional Modulation for Zero-shot Learning in Object Recognition <i>Now at PhD student at UCL. Co-supervised with Boris Katz.</i> | 2021 M9 |
| Christopher Z. Wang Modeling child language acquisition as grounded semantic parsing <i>Now at PhD student at MIT. Co-supervised with Boris Katz.</i> | 2021 M10 |
| Allison Fu Cognitively-plausible translation without aligned corpora <i>Co-supervised with Boris Katz.</i> | 2021 M11 |

| | |
|--|-------------|
| Julian Alverio Robots that understand language <i>Co-supervised with Boris Katz.</i> | 2021 M12 |
| David Mayo Understanding human and machine object recognition at scale <i>Co-supervised with Boris Katz.</i> | 2018 M13 |
| Dalitso Banda Removing dangerous visual stimuli from videos <i>Co-supervised with Boris Katz.</i> | 2018 M14 |
| Battushig Myanganbayar Inverse graphics for hand-pose reconstruction <i>Co-supervised with Boris Katz now at Apple.</i> | 2018 M15 |
| 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 M16 |
| Sergey Voronov Common Sense Reasoning Through Imagination <i>Co-supervised with Boris Katz. Now a PhD student at University of Ohio.</i> | 2016 M17 |
| Nicolas Rakover A uniform Representation for Visual Concepts <i>Co-supervised with Boris Katz. Now at Google.</i> | 2016 M18 |

Undergraduate students supervised

| | |
|--|----------------------|
| Nazar Ilamanov <i>Co-supervised with Boris Katz.</i> | <i>present</i> U1 |
| Julian Alverio <i>Co-supervised with Boris Katz. Now a Masters student at MIT.</i> | 2018 U2 |
| Gil Dekel <i>Co-supervised with Boris Katz. Now at Google.</i> | 2016-2018 U3 |
| Battushig Myanganbayar <i>Co-supervised with Boris Katz. Now a Masters student at MIT.</i> | 2016 U4 |
| David Mayo <i>Co-supervised with Boris Katz. Now a Masters student at MIT.</i> | 2016 U5 |
| Corinn Herrik <i>Co-supervised with Boris Katz.</i> | 2016 U6 |
| Sarah Coe <i>Co-supervised with Boris Katz.</i> | 2015-2016 U7 |
| Candace Ross <i>Co-supervised with Boris Katz. Now a PhD student at MIT.</i> | 2015 U8 |
| Victor Cabrera <i>Co-supervised with Boris Katz.</i> | 2015 U9 |
| Matthew Evanusa <i>Co-supervised with Boris Katz.</i> | 2014 U10 |
| Matthew Arbesfeld <i>Co-supervised with Boris Katz.</i> | 2014-2015 U11 |