Curriculum Vitae

July 19, 2025

Siddharth Narayanaswamy

School of Informatics
University of Edinburgh
10 Crichton St, Newington, Edinburgh EH8 9AB

www.iffsid.com
siddharth@iffsid.com
n.siddharth@ed.ac.uk
+44 7874 269532

Professional Experience

Reader in Explainable AI,	School of Informatics, University of Edinburgh	*July 2020 —*
Safe & Ethical AI Senior Research Fellow,	The Alan Turing Institute	*July 2022 — July 2024*
Safe & Ethical AI Research Fellow,	The Alan Turing Institute	*Nov 2019 — May 2021*
Senior Researcher,	Department of Engineering Science, University of Oxford	*May 2019 — July 2020*
Postdoctoral Research Assistant,	Department of Engineering Science, University of Oxford	*Jan 2016 — Apr 2019*
Postdoctoral Scholar,	Department of Psychology, Stanford University	*Apr 2014 — Dec 2015*
Graduate Research Assistant,	School of Electrical & Computer Engineering, Purdue University	*Aug 2008 — Mar 2014*

Education

Doctor of Philosophy (PhD), School of Electrical & Computer Engineering, Purdue University — *Aug 2008 — Mar 2014*
Thesis: Compositionality in Vision and Language

Bachelor of Engineering (BE), Electronics & Communication, Anna University, India — *Aug 2004 — Aug 2008*

Grants & Fellowships

AI Hub in Generative Models	Co-Lead (Methodology)	*Feb 2024 —*
Turing — Advanced Autonomy through Human-AI collaboration	Co-I	*Oct 2021 — Dec 2022*
Edinburgh Lab for Integrated AI	PI	*Oct 2021 — Apr 2024*
Huawei Edinburgh Lab	PI	*Jul 2021 — Jul 2025*
Facebook ParlAI Research Award	Co-Investigator	*Oct 2017 — Nov 2018*
EPSRC MURI Grant EP/N019474/1 — Commonsense Visual Reasoning	Senior/Key personnel	*Jan 2016 — Aug 2020*

Peer-Reviewed Journal Articles

J7 M. Proszewska, N. Malkin, and N. Siddharth.
On Designing Diffusion Autoencoders for Efficient Generation and Representation Learning
arXiv e-prints (May 2025). arXiv: 2506.00136 [stat.ML].

J6 R. de Bem, A. Ghosh, T. Ajanthan, O. Miksik, A. Boukhayma, N. Siddharth*, and P. H. S. Torr*.
DGPose: Deep Generative Models for Human Body Analysis
International Journal of Computer Vision (IJCV) 128.5 (2020), pp. 1537–1563.

J5 A. Muryy, N. Siddharth, N. Nardelli, A. Glennerster, and P. H. S. Torr.
Lessons from Reinforcement Learning for Biological Representations of Space
Vision Research (VR) 174 (2020), pp. 79–93.

J4 D. P. Barrett*, A. Barbu*, N. Siddharth*, and J. M. Siskind.
Saying What You're Looking For: Linguistics Meets Video Search
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 38.10 (2016), pp. 2069–2081.

J3 H. Yu*, N. Siddharth*, A. Barbu*, and J. M. Siskind.
A Compositional Framework for Grounding Language Inference, Generation, and Acquisition in Video
Journal of Artificial Intelligence Research (JAIR) 52 (2015), pp. 601–713.

J2 A. Barbu*, N. Siddharth*, A. Michaux, and J. M. Siskind.
Simultaneous Object Detection, Tracking, and Event Recognition
Advances in Cognitive Systems (ACS) 2 (2012), pp. 203–220.

J1 N. Siddharth*, A. Barbu*, and J. M. Siskind.
Seeing Unseeability to See the Unseeable
Advances in Cognitive Systems (ACS) 2 (2012), pp. 77–94.

Peer-Reviewed Conference Proceedings

C30 M. Opper and N. Siddharth.
Banyan: Improved Representation Learning with Explicit Structure
In *International Conference on Machine Learning (ICML)*. July 2025.

C29 A. B. Palmarini, C. G. Lucas, and N. Siddharth.
DreamDecompiler: Improved Bayesian Program Learning by Decompiling Amortised Knowledge
In *International Conference on Machine Learning (ICML)*. July 2024.

C28 V. Prokhorov, I. Titov, and N. Siddharth.
Autoencoding Conditional Neural Processes for Representation Learning
In *International Conference on Machine Learning (ICML)*. July 2024.

C27 C. Sun, Z. Yuan, K. Xu, L. Mai, N. Siddharth, S. Chen, and M. Marina.
Learning High-Frequency Functions Made Easy with Sinusoidal Positional Encoding
In *International Conference on Machine Learning (ICML)*. July 2024.

C26 W. Long, N. Siddharth, and B. Webber.
Multi-Label Classification for Implicit Discourse Relation Recognition
In *Findings of the ACL*. May 2024.

C25 M. Opper, V. Prokhorov, and N. Siddharth.
StrAE: Autoencoding for Pre-Trained Embeddings using Explicit Structure
In *Empirical Methods in Natural Language Processing (EMNLP)*. Dec. 2023.

C24 Y. Liang, J. Tenenbaum, T.-A. Le, and N. Siddharth.
Drawing out of Distribution with Neuro-Symbolic Generative Models
In *Advances in Neural Information Processing Systems (NeurIPS)*. Dec. 2022.

C23 Y. Shi, N. Siddharth, P. Torr, and A. R. Kosiorek.
Adversarial Masking for Self-Supervised Learning
In *International Conference on Machine Learning (ICML)*. June 2022.

C22 T. Joy, Y. Shi, P. H. S. Torr, T. Rainforth, S. M. Schmon, and N. Siddharth.
Learning multimodal VAEs through mutual supervision
In *International Conference on Learning Representations (ICLR)*. May 2022.

C21 T. A. Le, K. M. Collins, L. Hewitt, K. Ellis, N. Siddharth, S. J. Gershman, and J. B. Tenenbaum.
Hybrid Memoised Wake-Sleep: Approximate Inference at the Discrete-Continuous Interface
In *International Conference on Learning Representations (ICLR)*. May 2022.

C20 N. Miao, E. Mathieu, N. Siddharth, Y. W. Teh, and T. Rainforth.
On incorporating inductive biases into VAEs
In *International Conference on Learning Representations (ICLR)*. May 2022.

C19 Y. Shi, J. Seely, P. Torr, N. Siddharth, A. Hannun, N. Usunier, and G. Synnaeve.
Gradient Matching for Domain Generalization
In *International Conference on Learning Representations (ICLR)*. May 2022.

C18 T. Joy, S. Schmon, P. Torr, N. Siddharth*, and T. Rainforth*.
Capturing Label Characteristics in VAEs
In *International Conference on Learning Representations (ICLR)*. May 2021.

C17 Y. Shi, B. Paige, P. Torr, and N. Siddharth.
Relating by Contrasting: A Data-efficient Framework for Multimodal Generative Models
In *International Conference on Learning Representations (ICLR)*. May 2021.

C16 M. Igl, A. Gambardella, N. Nardelli, N. Siddharth, W. Böhmer, and S. Whiteson.
Multitask Soft Option Learning
In *Uncertainty in Artificial Intelligence (UAI)*. Aug. 2020.

C15 Y. Shi*, N. Siddharth*, B. Paige, and P. H. S. Torr.
Variational Mixture-of-Experts Autoencoders for Multi-Modal Deep Generative Models
In *Advances in Neural Information Processing Systems (NeurIPS)*. Dec. 2019, pp. 15692–15703.

C14 T. A. Le, A. R. Kosiorek, N. Siddharth, Y. W. Teh, and F. Wood.
Revisiting Reweighted Wake-Sleep for Models with Stochastic Control Flow
In *Uncertainty in Artificial Intelligence (UAI)*. July 2019.

C13 E. Mathieu*, T. Rainforth*, N. Siddharth*, and Y. W. Teh.
Disentangling Disentanglement in Variational Autoencoders
In *International Conference on Machine Learning (ICML)*. June 2019, pp. 4402–4412.

C12 B. Esmaeili, H. Wu, S. Jain, A. Bozkurt, N. Siddharth, B. Paige, D. H. Brooks, J. Dy, and J.-W. van de Meent.
Structured Disentangled Representations
In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. Apr. 2019, pp. 2525–2534.

C11 R. De Bem, A. Ghosh, A. Boukhayma, T. Ajanthan, N. Siddharth*, and P. Torr*.
A conditional deep generative model of people in natural images
In *Winter Conference on Applications of Computer Vision (WACV)*. Jan. 2019, pp. 1449–1458.

C10 S. Webb, A. Golinski, R. Zinkov, N. Siddharth, T. Rainforth, Y. W. Teh, and F. Wood.
Faithful Inversion of Generative Models for Effective Amortized Inference
In *Advances in Neural Information Processing Systems (NeurIPS)*. Dec. 2018, pp. 3074–3084.

C9 D. Massiceti, N. Siddharth, P. K. Dokania, and P. H. S. Torr.
FlipDial: A Generative Model for Two-Way Visual Dialogue
In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2018, pp. 6097–6105.

C8 N. Siddharth, B. Paige, J.-W. van de Meent, A. Desmaison, N. D. Goodman, P. Kohli, F. Wood, and P. H. S. Torr.
Learning Disentangled Representations with Semi-Supervised Deep Generative Models
In *Advances in Neural Information Processing Systems (NeurIPS)*. Dec. 2017, pp. 5927–5937.

C7 A. Barbu*, D. P. Barrett*, W. Chen, N. Siddharth*, C. Xiong, J. J. Corso, C. D. Fellbaum, S. J. H. Catherine Hanson, S. Helie, E. Malaia, B. A. Pearlmutter, J. M. Siskind, T. M. Talavage, and R. B. Wilbur.
Seeing is Worse than Believing: Reading People's Minds Better than Computer-Vision Methods Recognize Actions
In *Proceedings of the European Conference on Computer Vision (ECCV)*. Sept. 2014, pp. 612–627.

C6 N. Siddharth*, A. Barbu*, and J. M. Siskind.
Seeing What You're Told: Sentence-Guided Activity Recognition in Video
In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2014, pp. 732–739.

C5 Y. Cao, D. P. Barrett, A. Barbu, N. Siddharth, H. Yu, A. Michaux, Y. Lin, S. Dickinson, J. M. Siskind, and S. Wang.
Recognizing Human Activities from Partially Observed Videos
In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2013, pp. 2658–2665.

C4 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.
Video In Sentences Out
In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI)*. Aug. 2012, pp. 102–112.

C3 N. Siddharth*, A. Barbu*, and J. M. Siskind.
A Visual Language Model for Estimating Object Pose and Structure in a Generative Visual Domain
In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. May 2011, pp. 4854–4860.

C2 A. Barbu*, N. Siddharth*, and J. M. Siskind.
Learning Physically-Instantiated Game Play Through Visual Observation
In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. May 2010, pp. 1879–1886.

C1 N. Siddharth, M. Manivannan, S. Devasahayam, and G. Mathew.
Design of a Do-It-Yourself VR based Laparoscopic Simulator
In *Medicine Meets Virtual Reality (MMVR)*. 2009.

Technical Reports

T5 D. Massiceti, V. Kulharia, P. K. Dokania, N. Siddharth, and P. H. S. Torr.
A Revised Generative Evaluation of Visual Dialogue
ArXiv e-prints (Apr. 2020). arXiv: 2004.09272 [cs.CV].

T4 S. Bhatti, A. Desmaison, O. Miksik, N. Nardelli, N. Siddharth*, and P. H. S. Torr*.
Playing Doom with SLAM-Augmented Deep Reinforcement Learning
ArXiv e-prints (Dec. 2016). arXiv: 1612.00380 [cs.AI].

T3 A. Stuhlmüller, R. X. D. Hawkins, N. Siddharth, and N. D. Goodman.
Coarse-to-Fine Sequential Monte Carlo for Probabilistic Programs
ArXiv e-prints (Sept. 2015). arXiv: 1509.02962 [cs.AI].

T2 A. Barbu*, N. Siddharth*, C. Xiong, J. J. Corso, C. D. Fellbaum, C. Hanson, S. J. Hason, S. Helie, E. Malaia, B. A. Pearlmutter, J. M. Siskind, T. M. Talavage, and R. B. Wilbur.
The Compositional Nature of Verb and Argument Representations in the Human Brain
ArXiv e-prints (June 2013). arXiv: 1306.2293 [q-bio.NC].

T1 A. Barbu*, A. Bridge, D. Coroian, S. Dickinson, S. Mussman, N. Siddharth*, D. Salvi, L. Schmidt, J. Shangguan, J. M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang.
Large-Scale Automatic Labeling of Video Events with Verbs Based on Event-Participant Interaction
ArXiv e-prints (Apr. 2012). arXiv: 1204.3616 [cs.CV].

Workshop Publications

W15 M. Proszewska, N. Malkin, and N. Siddharth.
Bernoulli Priors as Efficient Denoising Guides for Diffusion Models
In *CVPR Workshop on Generative Models for Computer Vision (GMCV)*. June 2025.

W14 M. Jian and N. Siddharth.
Are LLMs good pragmatic speakers?
In *NeurIPS Workshop on Behavioural ML*. 2024.

W13 M. Proszewska and N. Siddharth.
Graph Kernel Convolutions for Interpretable Classification
In *ICLR Workshop on Data-Centric Machine Learning Research (DMLR)*. 2024.

W12 M. Opper, J. Morrison, and N. Siddharth.
On the Effect of Curriculum Learning with Developmental Data for Grammar Acquisition
In *EMNLP Workshop CoNLL-CMCL Shared Task BabyLM Challenge*. 2023.

W11 I. Vegner, N. Siddharth, and L. A. A. Doumas.
FoVAE: Reconstructive Foveation as a Self-Supervised Variational Inference Task for Visual Representation Learning
In *NeurIPS Workshop on Gaze Meets ML*. 2023.

W10 M. Opper, V. Prokhorov, and N. Siddharth.
StrAE: Autoencoding for Pre-Trained Embeddings using Explicit Structure
In *EMNLP Workshop on Unimodal & Multimodal Induction of Linguistic Structures*. 2022.

W9 N. Siddharth and B. Paige.
Learning Generative Models from Classifier Uncertainties
In *ICML Workshop on Uncertainty & Robustness in Deep Learning*. 2020.

W8 C. S. de Witt, B. Gram-Hansen, N. Nardelli, A. Gambardella, R. Zinkov, P. Dokania, N. Siddharth, A. B. Espinosa-Gonzalez, A. Darzi, and P. H. S. Torr.
Simulation-Based Inference for Global Health Decisions
In *ICML Workshop on ML for Global Health*. 2020.

W7 R. de Bem, A. Ghosh, T. Ajanthan, O. Miksik, N. Siddharth*, and P. H. S. Torr*.
A Semi-supervised Deep Generative Model for Human Body Analysis
In *ECCV Workshop on Human Behaviour Understanding*. 2018.

W6 M. Igl, W. Boehmer, A. Gambardella, P. H. S. Torr, N. Nardelli, N. Siddharth, and S. Whiteson.
Inference and Distillation for Option Learning
In *NeurIPS Workshop on Infer to Control*. 2018.

W5 D. Massiceti*, P. K. Dokania*, N. Siddharth*, and P. H. S. Torr.
Visual Dialogue without Vision or Dialogue
In *NeurIPS Workshop on Critiquing and Correcting Trends in Machine Learning*. 2018.

W4 E. Mathieu*, T. Rainforth*, N. Siddharth*, and Y. W. Teh.
Disentangling Disentanglement
In *NeurIPS Workshop on Bayesian Deep Learning*. 2018.

W3 N. Siddharth, B. Paige, J.-W. van de Meent, A. Desmaison, N. D. Goodman, P. Kohli, F. Wood, and P. H. S. Torr.
Inducing Interpretable Representations with Variational Autoencoders
In *NIPS Workshop on Interpretable Machine Learning in Complex Systems*. 2016.

W2 A. Barbu*, N. Siddharth*, and J. M. Siskind.
Language-Driven Video Retrieval
In *CVPR Workshop on Vision Meets Cognition*. 2014.

W1 N. Siddharth and N. D. Goodman.
Informative Scene Descriptions
In *CVPR Workshop on Language and Vision*. 2014.

Patents

P1 A. Barbu, N. Siddharth, H. Yu, and J. M. Siskind.
Correlating Videos and Sentences. U.S. pat. 9183466. Nov. 2015.

Invited Presentations

Efficient Modelling & Representation Learning	UCL AI Center	*Jun 2025*
Multimodal Representation Learning	ICLR 2023 workshop	*May 2023*
Bayesflow Research Talk	Google Brain — Bayesflow	*Jun 2022*
Hybrid Artificial Intelligence	AAAI	*Feb 2021*
CAIDA Research Seminar	University of British Columbia	*Dec 2019*
International Multimodal Communication Centre (IMCC) Seminar	University of Oxford	*Nov 2019*
Psychology Research Seminar	University of Glasgow	*Jan 2019*
Seminar on Joint Processing of Language and Visual Data	Dagstuhl	*Jan 2019*
Foundations of Situated and Multimodal Communication	IWCS, Montepellier	*Sept 2017*
Artificial Intelligence and Natural Computation Seminar	University of Birmingham	*May 2017*
ATI Probabilistic Programming Workshop	Alan Turing Institute	*Feb 2016*
Artificial Intelligence and Natural Computation Seminar	University of Birmingham	*Sept 2013*
COGS Research Seminar	University of Sussex	*Oct 2013*
Machine Learning and Perception Seminar	Microsoft Research (Cambridge)	*Oct 2013*

Professional Services

Area Chair	NeurIPS, ICML
PC Member	IJCAI 2016
Reviewing	AAAI, CVPR, ICRA, ICDL, ICLR, ICML, IJCAI, NeurIPS • PAMI, IJCV, PRL
Workshops	Language and Vision Workshop — CVPR 2015, 2017, 2018, 2019
	Learning Disentangled Representations — NIPS 2017

Teaching

IAML	Applied Machine Learning	School of Informatics, University of Edinburgh	*2022 —*
IAML	Introduction to Applied Machine Learning	School of Informatics, University of Edinburgh	*2021 — 2022*
DME	Data Mining and Exploration	School of Informatics, University of Edinburgh	*2021*
HDS-M02	Bayesian Modelling	Health Data Science (EPSRC CDT), University of Oxford	*2019 — 2019*
	Games Programming in C++	Coursework Module, Dept. of Engineering, University of Oxford	*2016 — 2018*
EE570	Artificial Intelligence	TA (Unofficial), Jeffrey Mark Siskind, Purdue University	*2008 — 2012*
EE473	Introduction to Artificial Intelligence	TA (Unofficial), Jeffrey Mark Siskind, Purdue University	*2008 — 2012*

Supervision

Mattia Opper	PhD [ongoing]
Magdalena Proszewska	PhD [ongoing]
Yuge Shi	DPhil 2023/Oxford [now at DeepMind]
Tom Joy	DPhil 2023/Oxford [now at Five AI / BOSCH]