DAVID COX, PH.D.

MIT-IBM Watson AI Lab ⋄ 314 Main St ⋄ Cambridge, MA 02142
(857) · 209 · 4089 ⋄ david.daniel.cox@gmail.com

EDUCATION

Massachusetts Institute of Technology *2007*
Ph.D. in Neuroscience (specialization in Computational Neuroscience)

Harvard Univeristy *2000*
A.B. in Biology and Psychology

EXPERIENCE

IBM Research March 2023 - Present
Vice President, AI Models *Cambridge, MA*

As VP for AI Models, I lead a global organization responsible for the development of large language models (LLMs) for language and code at IBM, along with supporting technologies (e.g. for efficient inference, trustworthy generative AI, etc).

IBM Research February 2018 - Present
IBM Director, MIT-IBM Watson AI Lab *Cambridge, MA*

The MIT-IBM Watson AI Lab is a first-of-its-kind academic-industrial partnership focused on advanced AI research, which was established with a $240m, 10 year commitment from IBM. As the Lab's inaugural IBM Director, I lead a team of researchers and engineers at IBM Research Cambridge who engage in a wide range of collaborative AI research projects with MIT faculty, students and staff. Through our MIT-IBM Lab Member program, we also engage with leading corporations across a wide range of industries to bridge from cutting edge science to real-world impact.

Harvard University 2007 - 2019

John L. Loeb Associate Professor of Natural Sciences and Engineering and Applied Sciences (2017 - 2019)
Assistant Professor of Molecular and Cellular Biology and of Computer Science (2012-2017)
Rowland Fellow (2007-2012)

As faculty in the Harvard Center for Brain Science, I led a laboratory with the mission of studying the biological underpinnings of intelligence and using those insights to design artificially intelligent systems. My experimental laboratory employed a range of different experimental techniques, from large scale microelectrode recordings and laser microscopy in living brains, to automated high-throughput behavioral assay and psychophysical experiments. Meanwhile, the computational arm of my lab designed biologically-inspired algorithms for computer vision and learning.

While at Harvard, I also created "The Fundamentals of Neuroscience," an online introductory neuroscience course that was one of the first Massive Open Online Courses (MOOCs) fielded by Harvard. To date, this course has enrolled well over half a million students from around the world and is the most widely taken neuroscience course ever.

DeepHealth 2017 - 2018
Co-founder *Cambridge, MA*

DeepHealth was founded in 2017 with the mission of building AI and computer vision systems to distill lifetimes of insights from medical experts into software to assist radiologists. DeepHealth was acquired in 2020 by RadNet, the largest outpatient radiology provider in North America, and received FDA approval for its first mammography product in 2021.

Perceptive Automata 2015 - 2018
Co-founder *Cambridge, MA*

Perceptive Automata was founded in 2015 with the goal of helping autonomous systems understand people—inferring their intentions, predicting their actions, and reasoning about what they know and don't know. For autonomous system (e.g. self-driving vehicles) understanding the humans around them is arguably the single biggest remaining problem for the widespread deployment. Perceptive has received over $20m in venture funding from variety of sources, including First Round capital and Toyota Ventures, and was named a World Economic Forum Tech Pioneer in 2019.

AWARDS, HONORS, AND FELLOWSHIPS

Roslyn Abramson Teaching Award – 2017

Google Faculty Research Award – 2010, 2014

Richard A. and Susan F. Smith Family Award for Excellence in Biomedical Research – 2013

Rowland Junior Fellowship, Harvard University – 2007-2012

National Defense Science and Engineering Graduate Fellowship – 2003-2006

Harvard Committee on Undergraduate Education, Certificate of Distinction in Teaching – 2000, 2001

PUBLICATIONS

Citations: 17515, h-index: 53, i10-index: 89
Google Scholar: https://scholar.google.com/citations?user=6S-WgLkAAAAJ

[1] Z. Sun, Y. Shen, Q. Zhou, H. Zhang, Z. Chen, D. Cox, Y. Yang, and C. Gan. Principle-driven self-alignment of language models from scratch with minimal human supervision. *Advances in Neural Information Processing Systems*, 36, 2024.

[2] P. Wang, R. Panda, L. T. Hennigen, P. Greengard, L. Karlinsky, R. Feris, D. D. Cox, Z. Wang, and Y. Kim. Learning to grow pretrained models for efficient transformer training. *arXiv preprint arXiv:2303.00980*, 2023.

[3] Z. Sun, Y. Shen, H. Zhang, Q. Zhou, Z. Chen, D. Cox, Y. Yang, and C. Gan. Salmon: Self-alignment with principle-following reward models. *arXiv preprint arXiv:2310.05910*, 2023.

[4] J. S. Smith, P. Cascante-Bonilla, A. Arbelle, D. Kim, R. Panda, D. Cox, D. Yang, Z. Kira, R. Feris, and L. Karlinsky. Construct-vl: Data-free continual structured vl concepts learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14994–15004, 2023.

[5] J. Y. Rhee, C. Echavarría, E. Soucy, J. Greenwood, J. A. Masís, and D. D. Cox. Neural correlates of visual object recognition in rats. *bioRxiv*, pages 2023–09, 2023.

[6] J. Masís, T. Chapman, J. Y. Rhee, D. D. Cox, and A. M. Saxe. Strategically managing learning during perceptual decision making. *Elife*, 12:e64978, 2023.

[7] L. Martie, J. Rosenberg, V. Demers, G. Zhang, O. Bhardwaj, J. Henning, A. Prasad, M. Stallone, J. Y. Lee, L. Yip, et al. Rapid development of compositional ai. *arXiv preprint arXiv:2302.05941*, 2023.

[8] C.-I. J. Lai, F. Shi, P. Peng, Y. Kim, K. Gimpel, S. Chang, Y.-S. Chuang, S. Bhati, D. Cox, D. Harwath, et al. Audio-visual neural syntax acquisition. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE, 2023.

[9] J. Kang, H. Luo, Y. Zhu, J. Glass, D. Cox, A. Ritter, R. Feris, and L. Karlinsky. Self-specialization: Uncovering latent expertise within large language models. *arXiv preprint arXiv:2310.00160*, 2023.

[10] M. Ding, Y. Xu, Z. Chen, D. D. Cox, P. Luo, J. B. Tenenbaum, and C. Gan. Embodied concept learner: Self-supervised learning of concepts and mapping through instruction following. In *Conference on Robot Learning*, pages 1743–1754. PMLR, 2023.

[11] Y. Zhang, W. Zhou, G. Zhang, D. Cox, and S. Chang. An adversarial framework for generating unseen images by activation maximization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3371–3379, 2022.

[12] J. Seale Smith, P. Cascante-Bonilla, A. Arbelle, D. Kim, R. Panda, D. Cox, D. Yang, Z. Kira, R. Feris, and L. Karlinsky. Construct-vl: Data-free continual structured vl concepts learning. *arXiv e-prints*, pages arXiv–2211, 2022.

[13] K. Qian, Y. Zhang, H. Gao, J. Ni, C.-I. Lai, D. Cox, M. Hasegawa-Johnson, and S. Chang. Contentvec: An improved self-supervised speech representation by disentangling speakers. In *International Conference on Machine Learning*, pages 18003–18017. PMLR, 2022.

[14] Y. Li, R. Panda, Y. Kim, C.-F. R. Chen, R. S. Feris, D. Cox, and N. Vasconcelos. Valhalla: Visual hallucination for machine translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5216–5226, 2022.

[15] C.-I. J. Lai, E. Cooper, Y. Zhang, S. Chang, K. Qian, Y.-L. Liao, Y.-S. Chuang, A. H. Liu, J. Yamagishi, D. Cox, et al. On the interplay between sparsity, naturalness, intelligibility, and prosody in speech synthesis. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8447–8451. IEEE, 2022.

[16] C.-I. Lai, F. Shi, P. Peng, Y. Kim, K. Gimpel, S. Chang, Y.-S. Chuang, S. Bhati, D. D. Cox, D. Harwath, et al. Textless phrase structure induction from visually-grounded speech. 2022.

[17] J. Dapello, K. Kar, M. Schrimpf, R. Geary, M. Ferguson, D. D. Cox, and J. J. DiCarlo. Aligning model and macaque inferior temporal cortex representations improves model-to-human behavioral alignment and adversarial robustness. *bioRxiv*, pages 2022–07, 2022.

[18] R. G. VidalMata, W. J. Scheirer, A. Kukleva, D. Cox, and H. Kuehne. Joint visual-temporal embedding for unsupervised learning of actions in untrimmed sequences. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1238–1247, 2021.

[19] K. Qian, Y. Zhang, S. Chang, J. Xiong, C. Gan, D. Cox, and M. Hasegawa-Johnson. Global prosody style transfer without text transcriptions. In *International Conference on Machine Learning*, pages 8650–8660. PMLR, 2021.

[20] C.-I. J. Lai, Y. Zhang, A. H. Liu, S. Chang, Y.-L. Liao, Y.-S. Chuang, K. Qian, S. Khurana, D. Cox, and J. Glass. Parp: Prune, adjust and re-prune for self-supervised speech recognition. *Advances in Neural Information Processing Systems (NeurIPS)*, 34, 2021.

[21] Y. Fu, Y. Zhang, Y. Zhang, D. Cox, and Y. Lin. Auto-nba: Efficient and effective search over the joint space of networks, bitwidths, and accelerators. In *International Conference on Machine Learning*, pages 3505–3517. PMLR, 2021.

[22] Y. Fu, Q. Yu, Y. Zhang, S. Wu, X. Ouyang, D. Cox, and Y. Lin. Drawing robust scratch tickets: Subnetworks with inborn robustness are found within randomly initialized networks. *Advances in Neural Information Processing Systems (NeurIPS*, 34, 2021.

[23] J. Dapello, J. Feather, H. Le, T. Marques, D. Cox, J. McDermott, J. J. DiCarlo, and S. Chung. Neural population geometry reveals the role of stochasticity in robust perception. *Advances in Neural Information Processing Systems (NeurIPS*, 34, 2021.

[24] G. Abrevaya, G. Dumas, A. Y. Aravkin, P. Zheng, J.-C. Gagnon-Audet, J. Kozloski, P. Polosecki, G. Lajoie, D. Cox, S. P. Dawson, et al. Learning brain dynamics with coupled low-dimensional nonlinear oscillators and deep recurrent networks. *Neural Computation*, 33(8):2087–2127, 2021.

[25] T. Schaffter, D. S. Buist, C. I. Lee, Y. Nikulin, D. Ribli, Y. Guan, W. Lotter, Z. Jie, H. Du, S. Wang, et al. Evaluation of combined artificial intelligence and radiologist assessment to interpret screening mammograms. *JAMA network open*, 3(3):e200265–e200265, 2020.

[26] K. Qian, Y. Zhang, S. Chang, M. Hasegawa-Johnson, and D. Cox. Unsupervised speech decomposition via triple information bottleneck. In *International Conference on Machine Learning*, pages 7836–7846. PMLR, 2020.

[27] J. Masís, T. Chapman, J. Y. Rhee, D. D. Cox, and A. M. Saxe. Rats strategically manage learning during perceptual decision making. *bioRxiv*, 2020.

[28] W. Lotter, G. Kreiman, and D. Cox. A neural network trained for prediction mimics diverse features of biological neurons and perception. *Nature Machine Intelligence*, 2(4):210–219, 2020.

[29] S. Han, A. Srivastava, C. L. Hurwitz, P. Sattigeri, and D. D. Cox. not-so-biggan: Generating high-fidelity images on a small compute budget. 2020.

[30] G. Guitchounts, J. Masis, S. B. Wolff, and D. Cox. Encoding of 3d head orienting movements in the primary visual cortex. *Neuron*, 108(3):512–525, 2020.

[31] G. Guitchounts, W. Lotter, J. Dapello, and D. Cox. Stable 3d head direction signals in the primary visual cortex. *bioRxiv*, 2020.

[32] G. Guitchounts and D. Cox. 64-channel carbon fiber electrode arrays for chronic electrophysiology. *Scientific Reports*, 10(1):1–9, 2020.

[33] C. Gan, J. Schwartz, S. Alter, M. Schrimpf, J. Traer, J. De Freitas, J. Kubilius, A. Bhandwaldar, N. Haber, M. Sano, et al. Threedworld: A platform for interactive multi-modal physical simulation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[34] J. Dapello, T. Marques, M. Schrimpf, F. Geiger, D. D. Cox, and J. J. DiCarlo. Simulating a primary visual cortex at the front of cnns improves robustness to image perturbations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[35] L. Amini, C.-H. Chen, D. Cox, A. Oliva, and A. Torralba. Experiences and insights for collaborative industry-academic research in artificial intelligence. *AI Magazine*, 41(1):70–81, 2020.

[36] P. Tillet, H.-T. Kung, and D. Cox. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, pages 10–19, 2019.

[37] A. M. Saxe, Y. Bansal, J. Dapello, M. Advani, A. Kolchinsky, B. D. Tracey, and D. D. Cox. On the information bottleneck theory of deep learning. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124020, 2019.

[38] C. Gan, H. Zhao, P. Chen, D. Cox, and A. Torralba. Self-supervised moving vehicle tracking with stereo sound. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7053–7062, 2019.

[39] S. L. Farhi, V. J. Parot, A. Grama, M. Yamagata, A. S. Abdelfattah, Y. Adam, S. Lou, J. J. Kim, R. E. Campbell, D. D. Cox, et al. Wide-area all-optical neurophysiology in acute brain slices. *Journal of Neuroscience*, 39(25):4889–4908, 2019.

[40] Q. Fan, C.-F. R. Chen, H. Kuehne, M. Pistoia, and D. Cox. More is less: Learning efficient video representations by big-little network and depthwise temporal aggregation. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.

[41] J. DeBoer, C. Haney, S. Z. Atiq, C. Smith, and D. Cox. Hands-on engagement online: using a randomised control trial to estimate the impact of an at-home lab kit on student attitudes and achievement in a mooc. *European Journal of Engineering Education*, 44(1-2):234–252, 2019.

[42] X. Chen, S. Liu, K. Xu, X. Li, X. Lin, M. Hong, and D. Cox. Zo-adamm: Zeroth-order adaptive momentum method for black-box optimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.

[43] E. Wu, K. Wu, D. Cox, and W. Lotter. Conditional infilling gans for data augmentation in mammogram classification. In *Image analysis for moving organ, breast, and thoracic images*, pages 98–106. Springer, Cham, 2018.

[44] H. Tang, M. Schrimpf, W. Lotter, C. Moerman, A. Paredes, J. O. Caro, W. Hardesty, D. Cox, and G. Kreiman. Recurrent computations for visual pattern completion. *Proceedings of the National Academy of Sciences*, 115(35):8835–8840, 2018.

[45] O. Skocek, T. Nöbauer, L. Weilguny, F. Martínez Traub, C. N. Xia, M. I. Molodtsov, A. Grama, M. Yamagata, D. Aharoni, D. D. Cox, et al. High-speed volumetric imaging of neuronal activity in freely moving rodents. *Nature methods*, 15(6):429–432, 2018.

[46] I. Puri and D. D. Cox. A system for accurate tracking and video recordings of rodent eye movements using convolutional neural networks for biomedical image segmentation. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 3590–3593. IEEE, 2018.

[47] V. J. Parot, S. L. Farhi, A. Grama, M. Yamagata, A. Abdelfattah, Y. Adam, S. Lou, J. J. Kim, R. E. Campbell, D. D. Cox, et al. Wide area profiling of neuronal function using hadamard microscopy. In *Optics and the Brain*, pages BW2C–3. Optical Society of America, 2018.

[48] J. Masís, D. Mankus, S. B. Wolff, G. Guitchounts, M. Joesch, and D. D. Cox. A micro-ct-based method for quantitative brain lesion characterization and electrode localization. *Scientific Reports*, 8(1):1–14, 2018.

[49] R. C. Fong, W. J. Scheirer, and D. D. Cox. Using human brain activity to guide machine learning. *Scientific Reports*, 8(1):1–10, 2018.

[50] P. Tillet, H. Kung, and D. Cox. Infomax-ica using hessian-free optimization. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2537–2541. IEEE, 2017.

[51] P. Tillet and D. Cox. Input-aware auto-tuning of compute-bound hpc kernels. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–12, 2017.

[52] W. Lotter, G. Sorensen, and D. Cox. A multi-scale cnn and curriculum learning strategy for mammogram classification. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pages 169–177. Springer, Cham, 2017.

[53] W. Lotter, G. Kreiman, and D. Cox. Deep predictive coding networks for video prediction and unsupervised learning. In *International Conference of Learning Representations (ICLR)*. arXiv preprint arXiv:1605.08104, 2017.

[54] L. Isik, W. Lotter, N. Crone, D. Cox, N. Kanwisher, W. Andreson, and G. Kreiman. Task dependent modulation before, during and after visually evoked responses in human intracranial recordings. *Journal of Vision*, 17(10):983–983, 2017.

[55] C. Haney, J. Qureshey, S. Z. Atiq, D. Cox, and J. DeBoer. Understanding lived experiences of students in a blended-learning neuroscience mooc, 2017.

[56] C.-Y. Tsai, A. Saxe, and D. Cox. Tensor switching networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, 2016.

[57] C.-Y. Tsai and D. Cox. Characterizing visual representations within convolutional neural networks: Toward a quantitative approach. In *ICML Workshop on Vis for Deep Learning*, 2016.

[58] R. Schalek, D. Lee, N. Kasthuri, A. Peleg, T. Jones, V. Kaynig, D. Haehn, H. Pfister, D. Cox, and J. W. Lichtman. Imaging a 1 mm 3 volume of rat cortex using a multibeam sem. *Microscopy and Microanalysis*, 22(S3):582–583, 2016.

[59] J. P. Papa, W. Scheirer, and D. D. Cox. Fine-tuning deep belief networks using harmony search. *Applied Soft Computing*, 46:875–885, 2016.

[60] C. L. Haney, J. DeBoer, S. Z. Atiq, and D. Cox. Comparing different learning activities in a global neuroscience mooc. In *2016 ASEE Annual Conference & Exposition*, 2016.

[61] S. Z. Atiq, C. L. Haney, J. DeBoer, and D. Cox. Understanding student experiences in a blended-learning mooc: a phenomenographic study. In *2016 ASEE Annual Conference & Exposition*, 2016.

[62] D. Zoccolan, D. D. Cox, and A. Benucci. What can simple brains teach us about how vision works. *Frontiers in neural circuits*, 9:51, 2015.

[63] C. Tsai and D. Cox. Are deep learning algorithms easily hackable, 2015.

[64] J. P. Papa, G. H. Rosa, A. N. Marana, W. Scheirer, and D. D. Cox. Model selection for discriminative restricted boltzmann machines through meta-heuristic techniques. *Journal of Computational Science*, 9:14–18, 2015.

[65] S. Lowry, N. Sünderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford. Visual place recognition: A survey. *IEEE Transactions on Robotics*, 32(1):1–19, 2015.

[66] W. Lotter, G. Kreiman, and D. Cox. Unsupervised learning of visual structure using predictive generative networks. *ICLR*, 2015.

[67] J. Bergstra, N. Pinto, and D. D. Cox. Skdata: data sets and algorithm evaluation protocols in python. *Computational Science & Discovery*, 8(1):014007, 2015.

[68] J. Bergstra, B. Komer, C. Eliasmith, D. Yamins, and D. D. Cox. Hyperopt: a python library for model selection and hyperparameter optimization. *Computational Science & Discovery*, 8(1):014008, 2015.

[69] S. Z. Atiq, X. Chen, D. D. Cox, and J. DeBoer. International stem classrooms: The experiences of students around the world using physical remote laboratory kits. In *2015 ASEE International Forum*, pages 19–23, 2015.

[70] E. Vig, M. Dorr, and D. Cox. Large-scale optimization of hierarchical features for saliency prediction in natural images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2798–2805, 2014.

[71] W. J. Scheirer, S. E. Anthony, K. Nakayama, and D. D. Cox. Perceptual annotation: Measuring human vision to improve computer vision. *IEEE transactions on pattern analysis and machine intelligence*, 36(8):1679–1686, 2014.

[72] M. Milford, E. Vig, W. Scheirer, and D. Cox. Vision-based simultaneous localization and mapping in changing outdoor environments. *Journal of Field Robotics*, 31(5):780–802, 2014.

[73] M. Milford, W. Scheirer, E. Vig, A. Glover, O. Baumann, J. Mattingley, and D. Cox. Condition-invariant, top-down visual place recognition. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5571–5577. IEEE, 2014.

[74] D. D. Cox and T. Dean. Neural networks and neuroscience-inspired computer vision. *Current Biology*, 24(18):R921–R929, 2014.

[75] D. D. Cox. Do we understand high-level vision? *Current opinion in neurobiology*, 25:187–193, 2014.

[76] G. Chiachia, A. X. Falcao, N. Pinto, A. Rocha, and D. Cox. Learning person-specific representations from faces in the wild. *IEEE Transactions on Information Forensics and Security*, 9(12):2089–2099, 2014.

[77] M. Milford, E. Vig, W. Scheirer, and D. Cox. Towards condition-invariant, top-down visual place recognition. In *Australasian Conference on Robotics and Automation*, pages 1–10, 2013.

[78] M. Meister and D. Cox. Rats maintain a binocular field centered on the horizon. *F1000Research*, 2, 2013.

[79] M. Günther, A. Costa-Pazo, C. Ding, E. Boutellaa, G. Chiachia, H. Zhang, M. de Assis Angeloni, V. Štruc, E. Khoury, E. Vazquez-Fernandez, et al. The 2013 face recognition evaluation in mobile environment. In *2013 International Conference on Biometrics (ICB)*, pages 1–7. IEEE, 2013.

[80] J. Bergstra, D. Yamins, D. D. Cox, et al. Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. In *Proceedings of the 12th Python in science conference*, volume 13, page 20, 2013.

[81] J. Bergstra, D. Yamins, and D. Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning*, pages 115–123. PMLR, 2013.

[82] E. Vig, M. Dorr, and D. D. Cox. Saliency-based selection of sparse descriptors for action recognition. In *2012 19th IEEE International Conference on Image Processing*, pages 1405–1408. IEEE, 2012.

[83] E. Vig, M. Dorr, and D. Cox. Space-variant descriptor sampling for action recognition based on saliency and eye movements. In *European conference on computer vision*, pages 84–97. Springer, Berlin, Heidelberg, 2012.

[84] N. Pinto and D. D. Cox. High-throughput-derived biologically-inspired features for unconstrained face recognition. *Image and Vision Computing*, 30(3):159–168, 2012.

[85] N. Pinto and D. D. Cox. Gpu metaprogramming: A case study in biologically inspired machine vision. In *GPU Computing Gems Jade Edition*, pages 457–471. Morgan Kaufmann, 2012.

[86] K. Marra, B. Graham, S. Carouso, and D. Cox. Implementation of a peltier-based cooling device for localized deep cortical deactivation during in vivo object recognition testing. In *APS March Meeting Abstracts*, volume 2012, pages Q39–003, 2012.

[87] G. Chiachia, N. Pinto, W. R. Schwartz, A. Rocha, A. X. Falcão, and D. D. Cox. Person-specific subspace analysis for unconstrained familiar face identification. In *BMVC*, pages 1–12, 2012.

[88] J. Bergstra, N. Pinto, and D. Cox. Machine learning for predictive auto-tuning with boosted regression trees. In *2012 Innovative Parallel Computing (InPar)*, pages 1–9. IEEE, 2012.

[89] N. Pinto, Y. Barhomi, D. D. Cox, and J. J. DiCarlo. Comparing state-of-the-art visual features on invariant object recognition tasks. In *2011 IEEE workshop on Applications of computer vision (WACV)*, pages 463–470. IEEE, 2011.

[90] D. Cox and N. Pinto. Beyond simple features: A large-scale feature search approach to unconstrained face recognition. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 8–15. IEEE, 2011.

[91] D. Zoccolan, B. J. Graham, and D. D. Cox. A self-calibrating, camera-based eye tracker for the recording of rodent eye movements. *Frontiers in neuroscience*, 4:193, 2010.

[92] V. Sriram, K. Tsoi, W. Luk, and D. Cox. Design-space exploration of biologically-inspired visual object recognition algorithms using cpus, gpus, and fpgas. *Many-core and reconfigurable supercomputing*, 2010.

[93] V. Sriram, D. Cox, K. H. Tsoi, and W. Luk. Towards an embedded biologically-inspired machine vision processor. In *2010 International Conference on Field-Programmable Technology*, pages 273–278. IEEE, 2010.

[94] N. Pinto and D. Cox. An evaluation of the invariance properties of a biologically-inspired system for unconstrained face recognition. In *International Conference on Bio-Inspired Models of Network, Information, and Computing Systems*, pages 505–518. Springer, Berlin, Heidelberg, 2010.

[95] D. Zoccolan, N. Oertelt, J. J. DiCarlo, and D. D. Cox. A rodent model for the study of invariant visual object recognition. *Proceedings of the National Academy of Sciences*, 106(21):8748–8753, 2009.

[96] N. Pinto, D. Doukhan, J. J. DiCarlo, and D. D. Cox. A high-throughput screening approach to discovering good forms of biologically inspired visual representation. *PLoS computational biology*, 5(11):e1000579, 2009.

[97] N. Pinto, J. J. DiCarlo, and D. D. Cox. How far can you get with a modern face recognition test set using only simple features? In *2009 IEEE conference on computer vision and pattern recognition*, pages 2591–2598. IEEE, 2009.

[98] L. R. Mujica-Parodi, H. H. Strey, B. Frederick, R. Savoy, D. Cox, Y. Botanov, D. Tolkunov, D. Rubin, and J. Weber. Chemosensory cues to conspecific emotional stress activate amygdala in humans. *PloS One*, 4(7):e6415, 2009.

[99] N. Li, D. D. Cox, D. Zoccolan, and J. J. DiCarlo. What response properties do individual neurons need to underlie position and clutter "invariant" object recognition? *Journal of Neurophysiology*, 102(1):360–376, 2009.

[100] N. Pinto, J. J. DiCarlo, and D. D. Cox. Establishing good benchmarks and baselines for face recognition. In *Workshop on Faces In Real-Life Images: Detection, Alignment, and Recognition*, 2008.

[101] N. Pinto, D. D. Cox, and J. J. DiCarlo. Why is real-world visual object recognition hard? *PLoS computational biology*, 4(1):e27, 2008.

[102] D. D. Cox, A. M. Papanastassiou, D. Oreper, B. B. Andken, and J. J. DiCarlo. High-resolution three-dimensional microelectrode brain mapping using stereo microfocal x-ray imaging. *Journal of neurophysiology*, 2008.

[103] D. D. Cox and J. J. DiCarlo. Does learned shape selectivity in inferior temporal cortex automatically generalize across retinal position? *Journal of Neuroscience*, 28(40):10045–10055, 2008.

[104] J. J. DiCarlo and D. D. Cox. Untangling invariant object recognition. *Trends in cognitive sciences*, 11(8):333–341, 2007.

[105] D. D. Cox. *Reverse engineering object recognition*. PhD thesis, Massachusetts Institute of Technology, 2007.

[106] B. Balas, D. Cox, and E. Conwell. The effect of real-world personal familiarity on the speed of face information processing. *PloS one*, 2(11):e1223, 2007.

[107] B. Balas, E. Conwell, and D. Cox. The effect of personal familiarity on the speed of face recognition. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 28, 2006.

[108] D. Zoccolan, D. D. Cox, and J. J. DiCarlo. Multiple object response normalization in monkey inferotemporal cortex. *Journal of Neuroscience*, 25(36):8150–8164, 2005.

[109] D. D. Cox, P. Meier, N. Oertelt, and J. J. DiCarlo. 'breaking'position-invariant object recognition. *Nature neuroscience*, 8(9):1145–1147, 2005.

[110] D. Cox, E. Meyers, and P. Sinha. Contextually evoked object-specific responses in human visual cortex. *Science*, 2004.

[111] D. D. Cox and R. L. Savoy. fmri "brain reading": detecting and classifying distributed patterns of fmri activity in human visual cortex. *Neuroimage*, 19(2):261–270, 2003.