

This is a section of [doi:10.7551/mitpress/10177.001.0001](https://doi.org/10.7551/mitpress/10177.001.0001)

Visual Cortex and Deep Networks

Learning Invariant Representations

By: Tomaso A. Poggio, Fabio Anselmi

Citation:

Visual Cortex and Deep Networks: Learning Invariant Representations

By: Tomaso A. Poggio, Fabio Anselmi

DOI: [10.7551/mitpress/10177.001.0001](https://doi.org/10.7551/mitpress/10177.001.0001)

ISBN (electronic): 9780262336710

Publisher: The MIT Press

Published: 2016

Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.



The MIT Press

References

- [1] F. Anselmi, L. Rosasco, and T. Poggio. On invariance and selectivity in representation learning. arXiv:1503.05938; Center for Brains, Minds and Machines, CBMM Memo No. 029, 2015.
- [2] K. Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics* 36 (1980): 193–202.
- [3] M. Riesenhuber and T. Poggio. Hierarchical models of object recognition in cortex. *Nature Neuroscience* 2 (1999): 1019–1025.
- [4] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. arXiv:1311.2901, 2013.
- [5] M. C. Potter. Meaning in visual search. *Science* 187, no. 4180 (1975): 965–966.
- [6] S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature* 381, no. 6582 (1996): 520–522.
- [7] Q. Liao, J. Z. Leibo, and T. Poggio. Learning invariant representations and applications to face verification. In *Proceedings of the 2013 Conference on Neural Information Processing Systems*.
- [8] F. Anselmi, J. Z. Leibo, L. Rosasco, J. Mutch, A. Tacchetti, and T. Poggio. Unsupervised learning of invariant representations in hierarchical architectures. arXiv:1311.4158, 2014; *Theoretical Computer Science*, in press, 2015.
- [9] F. Anselmi, J. Z. Leibo, L. Rosasco, J. Mutch, A. Tacchetti, and T. Poggio. Magic materials: A theory of deep hierarchical architectures for learning sensory representations. Center for Biological and Computational Learning, MIT, 2013.
- [10] F. Anselmi and T. Poggio. Representation learning in sensory cortex: A theory. Center for Brains, Minds and Machines, CBMM Memo No. 026, 2014.
- [11] T. Poggio, J. Mutch, F. Anselmi, A. Tacchetti, L. Rosasco, and J. Z. Leibo. Does invariant recognition predict tuning of neurons in sensory cortex?. Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT, Technical Report 019, 2013; Center for Biological and Computational Learning (CBCL), MIT, Paper No. 313, 2013.
- [12] F. Anselmi, L. Rosasco, T. Cheston, and T. Poggio. Deep convolutional networks are hierarchical kernel machines. Center for Brains, Minds and Machines, CBMM Memo No. 035, 2015.
- [13] J. Z. Leibo, Q. Liao, F. Anselmi, and T. Poggio. The invariance hypothesis implies domain-specific regions in visual cortex. bioRxiv:10.1101/004473, 2014.
- [14] T. Poggio, J. Mutch, and L. Isik. Computational role of eccentricity dependent cortical magnification. arXiv:1406.1770; Center for Brains, Minds and Machines, CBMM Memo No. 017, 2014.

- [15] J. Z. Leibo, F. Anselmi, J. Mutch, A. F. Ebiara, W. Freiwald, and T. Poggio. View-invariance and mirror-symmetric tuning in a model of the macaque face processing system. In *Proceedings of the 2013 Conference on Computational and Systems Neuroscience*, poster 1–54.
- [16] T. Poggio and S. Smale. The mathematics of learning: Dealing with data. *Notices of the American Mathematical Society* 50, no. 5 (2003): 537–544.
- [17] T. Lee and S. Soatto. Video-based descriptors for object recognition. *Image and Vision Computing* 29 (2012): 639–652.
- [18] P. Niyogi, F. Girosi, and T. Poggio. Incorporating prior information in machine learning by creating virtual examples. In *Proceedings of the IEEE* 86 (1998): 2196–2209.
- [19] Y. Abu-Mostafa. Hints and the VC dimension. *Neural Computation* 5 (1993): 278–288.
- [20] D. Hubel and T. Wiesel. Receptive fields: Binocular interaction and functional architecture in the cat’s visual cortex. *Journal of Physiology* 160, no. 1 (1962): 106.
- [21] S. Chikkerur and T. Poggio. Approximations in the HMAX model. Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT, Technical Report 021, 2011.
- [22] M. Kouh and T. Poggio. A canonical neural circuit for cortical nonlinear operations. *Neural Computation* 20, no. 6 (2008): 1427–1451.
- [23] E. Adelson and J. Bergen. Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America* 2, no. 2 (1985): 284–299.
- [24] D. Gabor. Uncertainty principles and signal recovery. *Journal of the Institution of Electrical Engineers. Part III: Radio and Communication Engineering* 93, no. 26 (1946): 429–441.
- [25] D. L. Donoho and S. P. Stark. Uncertainty principles and signal recovery. *SIAM Journal on Applied Mathematics* 49, no. 3 (1989): 906–931.
- [26] D. Ringach. Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology* 88, no. 1 (2002): 455–463.
- [27] J. P. Jones and L. A. Palmer. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology* 58, no. 6 (1987): 1233–1258.
- [28] C. Niell and M. Stryker. Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience* 28, no. 30 (2008): 7520–7536.
- [29] H. Bülthoff and S. Edelman. Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences* 89, no. 1 (1992): 60–64.
- [30] M. J. Tarr. Rotating objects to recognize them: A case study on the role of viewpoint dependency in the recognition of three dimensional objects. *Psychonomic Bulletin and Review* 2, no. 1 (1995): 55–82.
- [31] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, no. 6583 (1996): 607–609.
- [32] S. Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics* 65, no. 10 (2012): 1331–1398.
- [33] M. Riesenhuber and T. Poggio. Models of object recognition. *Nature Neuroscience* 3 (2000): 1199–1204.
- [34] Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation* 1, no. 4 (1989): 541–551.
- [35] J. Z. Leibo, J. Mutch, and T. Poggio. Why the brain separates face recognition from object recognition. In *Proceedings of the 2011 Conference on Neural Information Processing Systems*.
- [36] D. Arathorn. Computation in the higher visual cortices: Map-seeking circuit theory and application to machine vision. In *Proceedings of the 2004 International Symposium on Information Theory*, 73–78.

- [37] L. Sifre, S. Mallat, and P. France. Combined scattering for rotation invariant texture analysis. In *Proceedings of the 2012 European Symposium on Artificial Neural Networks*, 127–132.
- [38] T. Serre, M. Kouh, C. Cadieu, U. Knoblich, G. Kreiman, and T. Poggio. A theory of object recognition: Computations and circuits in the feedforward path of the ventral stream in primate visual cortex. Center for Biological and Computational Learning (CBCL), MIT, Paper No. 259, 2005.
- [39] S. S. Chikkerur, T. Serre, C. Tan, and T. Poggio. What and where: A Bayesian inference theory of attention. *Vision Research* 50 (2010): 2233–2247.
- [40] D. George and J. Hawkins. A hierarchical Bayesian model of invariant pattern recognition in the visual cortex. In *Proceedings of the 2005 IEEE International Joint Conference on Neural Networks*, 1812–1817.
- [41] S. Geman. Invariance and selectivity in the ventral visual pathway. *Journal of Physiology (Paris)* 100, no. 4 (2006): 212–224.
- [42] S. Hochstein and M. Ahissar. View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron* 36, no. 5 (2002): 791–804.
- [43] H. Schulz-Mirbach. Constructing invariant features by averaging techniques. In *Proceedings of the 1994 International Conference on Pattern Recognition. Conference B*, 387–390.
- [44] B. Haasdonk and H. Burkhardt. Invariant kernel functions for pattern analysis and machine learning. *Machine Learning* 68 (July 2007): 35–61.
- [45] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G.R.G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research* 11 (2010): 1517–1561.
- [46] R. Rosenholtz. General-purpose localization of textured image regions. In *Proceedings of the 1999 SPIE Conference on Human Vision and Electronic Imaging*, 4: 454–460.
- [47] M. Riesenhuber and T. Poggio. Hierarchical models of object recognition. *Nature Neuroscience* 2 (1999): 1019–1025.
- [48] J. Freeman and E. Simoncelli. Metamers of the ventral stream. *Nature Neuroscience* 14 (2011): 1195–1201.
- [49] P. Földiák. Learning invariance from transformation sequences. *Neural Computation* 3, no. 2 (1991): 194–200.
- [50] T. Serre, A. Oliva, and T. Poggio. A feedforward architecture accounts for rapid categorization. *Proceedings of the National Academy of Sciences* 104, no. 15 (2007): 6424–6429.
- [51] Y. LeCun and Y. Bengio. Convolutional networks for images, speech, and time series. In *The Handbook of Brain Theory and Neural Networks*, ed. M. A. Arbib, 276–278. 2nd ed. Cambridge, MA: MIT Press, 2002.
- [52] T. Poggio. A theory of how the brain might work. *Cold Spring Harbor Symposia on Quantitative Biology* 55 (1990): 899–910.
- [53] D. Perrett and M. Oram. Neurophysiology of shape processing. *Image and Vision Computing* 11, no. 6 (1993): 317–333.
- [54] B. W. Mel. SEEMORE: Combining color, shape, and texture histogramming in a neurally inspired approach to visual object recognition. *Neural Computation* 9 (1997): 777–804.
- [55] S. M. Stringer and E. T. Rolls. Invariant object recognition in the visual system with novel views of 3D objects. *Neural Computation* 14 (2002): 2585–2596.
- [56] N. Pinto, J. J. DiCarlo, and D. Cox. How far can you get with a modern face recognition test set using only simple features? In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2591–2598.
- [57] A. M. Saxe, M. Bhand, R. Mudur, B. Suresh, and A. Y. Ng. Unsupervised learning models of primary cortical receptive fields and receptive field plasticity. In *Proceedings of the 2011 Conference on Neural Information Processing Systems*, 1971–1979.

- [58] Q. V. Le, R. Monga, M. Devin, G. Corrado, K. Chen, M. Ranzato, J. Dean, and A. Y. Ng. Building high-level features using large-scale unsupervised learning. arXiv:1112.6209, 2011.
- [59] O. Abdel-Hamid, A. Mohamed, H. Jiang, and G. Penn. Applying convolutional neural networks concepts to hybrid NN-HMM model for speech recognition. In *Proceedings of the 2012 IEEE International Conference on Acoustics, Speech, and Signal Processing*, 4277–4280.
- [60] D. O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley, 1949; new edition published 2002 by Lawrence Erlbaum Associates, Mahwah, NJ.
- [61] G. G. Turrigiano and S. B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience* 5, no. 2 (2004): 97–107.
- [62] E. Oja. Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology* 15, no. 3 (1982): 267–273.
- [63] J. Karhunen. Stability of Oja’s PCA subspace rule. *Neural Computation* 6 (1994): 739–747.
- [64] T. Sanger. Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural Networks* 2, no. 6 (1989): 459–473.
- [65] E. Oja. Principal components, minor components, and linear neural networks. *Neural Networks* 5, no. 6 (1992): 927–935.
- [66] A. Hyvriin and E. Oja. Independent component analysis by general nonlinear Hebbian-like learning rules. *Signal Processing* 64 (1998): 301–313.
- [67] D. Ruderman. The statistics of natural images. *Network: Computation in Neural Systems* 5 (1994): 517–548.
- [68] A. Torralba and A. Oliva. Statistics of natural image categories. *Network: Computation in Neural Systems* 14 (2003): 391–412.
- [69] J. Gallant, J. Braun, and D. V. Essen. Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual cortex. *Science* 259 (1993): 100–103.
- [70] J. Gallant, C. Connor, S. Rakshit, J. Lewis, and D. Van Essen. Neural responses to polar, hyperbolic, and Cartesian gratings in area V4 of the macaque monkey. *Journal of Neurophysiology* 76 (1996): 2718–2739.
- [71] J. Hegde and D. Van Essen. Selectivity for complex shapes in primate visual area V2. *Journal of Neuroscience* 20, no. 5 (2000): 61.
- [72] S. Mallat. *A Wavelet Tour of Signal Processing: The Sparse Way*. 3rd ed. Burlington, MA: Academic Press, 2008.
- [73] D. Marr, T. Poggio, and E. Hildreth. Smallest channel in early human vision. *Journal of the Optical Society of America* 70, no. 7 (1980): 868–870.
- [74] A. Cowey and E. T. Rolls. Human cortical magnification factor and its relation to visual acuity. *Experimental Brain Research* 21, no. 5 (1974): 447–454.
- [75] H. Strasburger, I. Rentschler, and M. Jüttner. Peripheral vision and pattern recognition: A review. *Journal of Vision* 11, no. 5 (2011): 13.
- [76] S. Ullmon. Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences* 11 (2): 58–69, 2007.
- [77] J. A. Mazer, W. E. Vinje, J. McDermott, P. H. Schiller, and J. L. Gallant. Spatial frequency and orientation tuning dynamics in area V1. *Proceedings of the National Academy of Sciences* 99, no. 3 (2002): 1645–1650.
- [78] D. Hubel and T. Wiesel. Uniformity of monkey striate cortex: A parallel relationship between field size, scatter, and magnification factor. *Journal of Comparative Neurology* 158 (1974): 295–305.
- [79] R. Gattass, C. Gross, and J. Sandell. Visual topography of V2 in the macaque. *Journal of Comparative Neurology* 201, no. 4 (1981): 519–539.

- [80] R. Gattass, A. Sousa, and C. Gross. Visuotopic organization and extent of V3 and V4 of the macaque. *Journal of Neuroscience* 8, no. 6 (1988): 1831–1845.
- [81] D. Levi and T. Carney. The effect of flankers on three tasks in central, peripheral, and amblyopic vision. *Journal of Vision* 11, no. 1 (2011): 10.
- [82] H. Bouma. Interaction effects in parafoveal letter recognition. *Nature* 226 (1970): 177–178.
- [83] D. Pelli and K. Tillman. The uncrowded window of object recognition. *Nature Neuroscience* 11, no. 10 (2008): 1129–1135.
- [84] D. Levi. Crowding, an essential bottleneck for object recognition: A minireview. *Vision Research* 48, no. 5 (2008): 635–654.
- [85] M. Lev, O. Yehezkel, and U. Polat. Uncovering foveal crowding? *Scientific Reports* 4 (2014): article 4067.
- [86] P. G. Schyns and F. Gosselin. Diagnostic use of scale information for componential and holistic recognition. In *Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes*, ed. M. A. Peterson and G. Rhodes, 120–145. New York: Oxford University Press, 2003.
- [87] S. Anstis. A chart demonstrating variations in acuity with retinal position. *Vision Research* 14, no. 7 (1974): 589–592.
- [88] H. D. Beeck and R. Vogels. Spatial sensitivity of macaque inferior temporal neurons. *Journal of Comparative Neurology* 426, no. 4 (2000): 505–518.
- [89] J. DiCarlo and J. Maunsell. Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *Journal of Neurophysiology* 89, no. 6 (2003): 3264–3278.
- [90] B. Balas, L. Nakano, and R. Rosenholtz. A summary-statistic representation in peripheral vision explains visual crowding. *Journal of Vision* 9, no. 12 (2009): 13.
- [91] D. Tsao and W. Freiwald. Faces and objects in macaque cerebral cortex. *Nature* 6, no. 9 (2003): 989–995.
- [92] J. Z. Leibo, J. Mutch, and T. Poggio. How can cells in the anterior medial face patch be viewpoint invariant? Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT, Technical Report 057, 2010.
- [93] N. Kanwisher. Functional specificity in the human brain: A window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences* 25, no. 107 (2010): 11163–11170.
- [94] P. Downing and Y. Jiang. A cortical area selective for visual processing of the human body. *Science* 293, no. 5539 (2001): 2470–2073.
- [95] M. Bar, E. Aminoff, and D. L. Schacter. Scenes unseen: The parahippocampal cortex intrinsically subserves contextual associations, not scenes or places per se. *Journal of Neuroscience* 28, no. 34: 8539–8544.
- [96] N. Logothetis, J. Pauls, H. Bülthoff, and T. Poggio. View-dependent object recognition by monkeys. *Current Biology* 4, no. 5. (1994): 401–414.
- [97] N. Logothetis, J. Pauls, and T. Poggio. Shape representation in the inferior temporal cortex of monkeys. *Current Biology* 5, no. 5 (1995): 552–563.
- [98] T. Poggio and S. Edelman. A network that learns to recognize three-dimensional objects. *Nature* 343 (1990): 263–266.
- [99] S. Ku, A. Tolias, N. Logothetis, and J. Goense. fMRI of the face-processing network in the ventral temporal lobe of awake and anesthetized macaques. *Neuron* 70, no. 2 (2011): 352–362.
- [100] R. Malach, J. B. Reppas, R. R. Benson, K. Kwong, H. Jiang, W. A. Kennedy, P. J. Ledden, T. J. Brady, B. R. Rosen, and R. B. Tootell. Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences* 92, no. 18 (1995): 8135–8139.

- [101] L. Cohen, S. Dehaene, L. Naccache, S. Lehéricy, G. Dehaene-Lambertz, M. Hénaff, and F. Michel. The visual word form area. *Brain* 123, no. 2 (2000): 291–307.
- [102] N. Kriegeskorte. Representational similarity analysis: Connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience* 2, no. 4 (2008).
- [103] N. Logothetis and D. Sheinberg. Visual object recognition. *Annual Review of Neuroscience* 19, no. 1 (1996): 577–621.
- [104] M. Stryker. Temporal associations. *Nature* 354 (1991): 108–109.
- [105] W. Freiwald and D. Tsao. Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* 330, no. 6005 (2010): 845–851.
- [106] W. Freiwald. A face feature space in the macaque temporal lobe. *Nature Neuroscience* 12 (2009): 1187–1196.
- [107] J. Z. Leibo. The invariance hypothesis and the ventral stream. Ph.D. diss., Massachusetts Institute of Technology, 2013.
- [108] J. Mutch, J. Z. Leibo, S. Smale, L. Rosasco, and T. Poggio. Neurons that confuse mirror-symmetric object views. Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT, Technical Report 062, 2010; Center for Biological and Computational Learning (CBCL), MIT, Paper No. 295, 2010.
- [109] Y. LeCun. LeNet-5, Convolutional Neural Networks. <http://yann.lecun.com/exdb/lenet/index.html>. Accessed Nov. 2015.
- [110] B. C. Motter and G. F. Poggio. Blmocular fixation in the rhesus monkey: spatial and temporal characteristics. *Exp Brain Res.* 54, 304–314, 1984.
- [111] E. Bricolo. On the representation of novel objects: Human psychophysics, monkey physiology, and computational models. Ph.D. diss., Massachusetts Institute of Technology, 1996.
- [112] D. Tsao and M. Livingstone. Mechanisms of face perception. *Annual Review of Neuroscience* 31 (2008): 411–437.
- [113] C. P. Hung, G. Kreiman, T. Poggio, and J. J. DiCarlo. Fast readout of object identity from macaque inferior temporal cortex. *Science* 310 (2005): 863–866.
- [114] A. Hyvriinen and E. Oja. Independent component analysis: Algorithms and applications. *Neural Networks* 13 (2000): 411–430.
- [115] E. Borenstein and S. Ullman. Combined top-down/bottom-up segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30, no. 12 (2008): 2109–2125.
- [116] D. Heeger. Normalization of cell responses in cat striate cortex. *Visual Neuroscience* 9 (1992):181–197.
- [117] T. Poggio and W. Reichardt. Considerations on models of movement detection. *Kybernetik* 13 (1973): 223–227.
- [118] J. Mutch and D. Lowe. Multiclass object recognition with sparse localized features. In *Proceedings of the 2006 IEEE Conference on Computer Vision and Pattern Recognition*, 11–18.
- [119] Y. LeCun, F. Huang, and L. Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *Proceedings of the 2004 IEEE Conference on Computer Vision and Pattern Recognition*, 2: 97–104.
- [120] G. Palm. On associative memory. *Biological Cybernetics* 36, no. 1 (1980): 19–31.
- [121] T. Poggio. On optimal nonlinear associative recall. *Biological Cybernetics* 19, no. 4 (1975): 201–209.
- [122] F. Cucker and S. Smale. On the mathematical foundations of learning. *Bulletin of the American Mathematical Society* 39, no. 1 (2001): 1–49.
- [123] H. Cramér and H. Wold. Some theorems on distribution functions. *Journal of the London Mathematical Society* 4 (1936): 290–294.
- [124] J. Cuesta-Albertos, R. Fraiman, and T. Ransford. A sharp form of the Cramér-Wold theorem. *Journal of Theoretical Probability* 20, no. 2 (2007): 201–209.

- [125] J. Cuesta-Albertos. How many random projections suffice to determine a probability distribution? IPMs sections, 2009.
- [126] A. Heppes. On the determination of probability distributions of more dimensions by their projections. *Acta Mathematica Hungarica* 7, no. 3 (1956): 403–410.
- [127] T. Plate. Holographic reduced representations: Convolution algebra for compositional distributed representations. In *Proceedings of the 1991 International Joint Conference on Artificial Intelligence*, 30–35.
- [128] L. Isik, E. M. Meyers, J. Z. Leibo, and T. Poggio. The timing of invariant object recognition in the human visual system. *Journal of Neurophysiology* 111, no. 1 (2013): 91–102.
- [129] J. Antoine, R. Murenzi, P. Vandergheynst, and S. Ali. *Two-Dimensional Wavelets and Their Relatives*. Cambridge: Cambridge University Press, 2004.
- [130] C. F. Stevens. Preserving properties of object shape by computations in primary visual cortex. *Proceedings of the National Academy of Sciences* 101, no. 11 (2004): 15524–15529.
- [131] A. Grossman, J. Morlet, and T. Paul. Transforms associated to square integrable group representations: Examples. *Annales de l'I.H.P. Probabilités et statistiques* 45 (1986): 293–309.
- [132] A. Saxe, P. W. Koh, Z. Chen, M. Bhand, B. Suresh, and A. Ng. On random weights and unsupervised feature learning. In *Proceedings of the 2011 International Conference on Machine Learning*, 1089–1096.
- [133] J. Nelson. Johnson-Lindenstrauss notes. http://web.mit.edu/minilek/www/jl_notes.pdf. Accessed Nov. 2015.
- [134] Center for Brains, Minds and Machines. CBMM Turing++ questions. <http://cbmm.mit.edu/research/cbmm-challenge>. Accessed Nov. 2015.
- [135] H.J.W. Müller-Kirsten. *Introduction to Quantum Mechanics*. Singapore: World Scientific, 2006.

© 2016 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.

Subject to such license, all rights are reserved.



Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.

Library of Congress Cataloging-in-Publication Data

Names: Poggio, Tomaso, author. | Anselmi, Fabio, author.

Title: Visual cortex and deep networks : learning invariant representations /
Tomaso A. Poggio and Fabio Anselmi.

Description: Cambridge, MA : MIT Press, [2016] | Series: Computational
neuroscience | Includes bibliographical references and index.

Identifiers: LCCN 2016005774 | ISBN 9780262034722 (hardcover : alk. paper)

Subjects: LCSH: Visual cortex. | Vision. | Neural networks (Neurobiology) |
Perceptual learning. | Computational neuroscience.

Classification: LCC QP383.15 .P64 2016 | DDC 612.8—dc23 LC record available at
<http://lccn.loc.gov/2016005774>

10 9 8 7 6 5 4 3 2 1