

参考文献

- [1] Goodfellow, I., Bengio, Y., and Courville, A. Deep Learning [M]. Cambridge: MIT Press, 2016.
- [2] Tom Mitchell, Machine Learning [M]. 1st ed. New York: McGraw Hill, 1997.
- [3] Wei-Lun Chao, Machine Learning Tutorial [OL]. Taiwan: DISP Lab, National Taiwan University, 2016. <http://disp.ee.ntu.edu.tw/~pujols/Machine Learning Tutorial.pdf>
- [4] Fei-Fei Li, Andrej Karpathy, Justin Johoson. CS231n: Convolutional Neural Networks for Visual Recognition [OL]. Palo Alto: Stanford University, 2015. <http://cs231n.stanford.edu>
- [5] Simon J. D. Prince. Computer Vision: Models, Learning and Inference [M]. Cambridge: Cambridge University Press, 2012.
- [6] Richard Szeliski, Computer Vision Algorithms and Applications [M]. New York: Cambridge University Press, 2010.
- [7] Y. Bengio, A. Courville and P. Vincent, Representation Learning: A Review and New Perspectives [J], IEEE Transactions on Pattern Analysis and Machine Intelligence, 2013, vol. 35, no. 8, pp. 1798-1828.
- [8] Sudha Bhingardive, Rudramurthy V, Kevin Patel, Prerana Singhal. Deep Learning and Distributed Representations [C]. 12th International Conference on Natural Language Processing, Trivandrum: Indian Institute of Information Technology and Management-Kerala, 2015.
- [9] Hugo Larochelle, Yoshua Bengio, Jérôme Louradour, and Pascal Lamblin. Exploring Strategies for Training Deep Neural Networks [J]. The Journal of Machine Learning Research, vol. 10, pp. 1-40, 2009.
- [10] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition [C]. Proceeding of International Conference on Machine Learning, Cambridge: MIT Press, 2014.
- [11] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition [J]. Proceeding of IEEE, New York: IEEE Press, 1998.
- [12] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks[C]. Advances in Neural Information Processing Systems, California: Neural Information Processing Systems Foundation, 2012.
- [13] 黄凯奇, 任伟强, 谭铁牛, 图像物体分类与检测算法综述 [J], 计算机学报, 2014 年 6 月, 第 37 卷, 第 6 期: 1225-1240
- [14] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding [C]. Proceedings of the 22nd ACM international conference on Multimedia: 675-678, 2014.
- [15] NVIDIA, Deep Learning on GPUs [OL]. NVIDIA, March 2016.

- http://www.nvidia.com/content/events/geoInt2015/LBrown_DL.pdf
- [16] Y. Lecun, Y. Bengio, G. Hinton, Deep Learning [J], Nature, vol. 521, no. 7553: 436-444, 2015.
- [17] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks [C]. Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition: 1725-1732, District of Columbia, Washington: IEEE Computer Society Washington, 2014.
- [18] Neocognitron [OL], Wikipedia. <https://en.wikipedia.org/wiki/Neocognitron>
- [19] J. Schmidhuber. Deep Learning in Neural Networks: An Overview [J]. Neural Networks, Volume 61: Pages 85-117, Amsterdam: Elsevier, 2015.
- [20] S. Bahrampour, N. Ramakrishnan, L. Schott, and M. Shah. Comparative study of caffe, neon, theano, and torch for deep learning [OL]. arXiv preprint, abs/1511.06435, 2015. <https://arxiv.org/abs/1511.06435>
- [21] F. Bastien, P. Lamblin, R. Pascanu, J. Bergstra, I. J. Goodfellow, A. Bergeron, N. Bouchard, and Y. Bengio. Theano: new features and speed improvements [C]. Deep Learning and Unsupervised Feature Learning workshop in The Conference of Advances in Neural Information Processing Systems, California: Neural Information Processing Systems Foundation.
- [22] Karl Tate, History of A.I.: Artificial Intelligence [OL]. Infographics Artist. Aug 25, 2014. <http://www.livescience.com/47544-history-of-a-i-artificial-intelligence-infographic.html>
- [23] Andrey Kurenkov. A 'Brief' History of Neural Nets and Deep Learning [OL], Dec 24, 2015. <http://www.andreykurenkov.com/writing/a-brief-history-of-neural-nets-and-deep-learning/>
- [24] 尼克, 人工智能的缘起 [J], 中国计算机学会通讯, 2016 年第三期《专栏》
- [25] David Auerbach, AI has grown up and left home [OL], <http://nautil.us/issue/8/home/ai-has-grown-up-and-left-home>
- [26] The 3rd NFN "Cognitive Vision" Workshop with the topic "History of Computer Vision" [OL], <http://www.icg.tugraz.at/News/historyOfCV>
- [27] Michael Nielsen. Neural Networks and Deep Learning [OL], <http://neuralnetworksanddeeplearning.com/>
- [28] Sheldon Axler. Linear Algebra Done Right [M]. 3rd ed. Berlin: Springer, 2015.
- [29] Fletcher Dunn, Ian Parberry. 3D Math Primer for Graphics and Game Development [M]. Bolingbrook : Wordware Publishing, Inc. 2002
- [30] Igor R. Shafarevich, Alexey Remizov, Linear Algebra and Geometry [M]. Berlin: Springer, 2013.
- [31] Why is log likelihood so widely used in Machine Learning [OL], <https://www.quora.com/Why-is-log-likelihood-so-widely-used-in-Machine-Learning>
- [32] Remi Emonet. Introduction to Sparsity in Modeling and Learning [OL]. University Jean Monnet - Laboratory Hubert Curien,

- <https://laboratoirehubertcurien.univ-st-etienne.fr/en/index.html>
- [33] Simon Foucart, Holger Rauhut. A Mathematical Introduction to Compressive Sensing. Applied and Numerical Harmonic Analysis [M]. Switzerland: Birkhauser, 2013.
 - [34] Yoshua Bengio. Curse of Dimensionality: Local versus Non-Local Learning [C], Canadian Institute for Advanced Research, Manifold Learning Workshop, 2005.
 - [35] Stephen Boyd and Lieven Vandenberghe, Convex Optimization [M]. Cambridge: Cambridge University Press, 2004.
 - [36] S. Becker and Y. LeCun, Improving the convergence of back-propagation learning with second-order methods [C], Proc. of the 1988 Connectionist Models Summer School, pages 29-37, San Mateo, 1988.
 - [37] Tom Schaul, Sixin Zhang, and Yann LeCun, No more pesky learning rates [C]. Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28 (ICML'13), Sanjoy Dasgupta and David McAllester (Eds.), 2013.
 - [38] Sargur Srihari. Machine Learning Tutorial, Chapter 5.4. The Hessian Matrix [OL], <http://www.cedar.buffalo.edu/~srihari/CSE574/Chap5/Chap5.4-Hessian.pdf>
 - [39] Matthew D Zeiler. Adadelta: an adaptive learning rate method [OL], arXiv preprint arXiv:1212.5701. 2012, <https://arxiv.org/abs/1212.5701>
 - [40] Brian Heinold. An Intuitive Guid to Numerical Methods [OL]. Department of Mathematics and Computer Science, Mount St. Mary's University, http://www.brianheinold.net/notes/An_Intuitive_Guide_to_Numerical_Methods_Heinold.pdf
 - [41] Dimitri P. Bertsekas and John N. Tsitsiklis. Introduction to Probability [M], 2nd ed. Nashua : Athena Scientific, 2004.
 - [42] Charles M. Grinstead, J. Laurie Snell. Introduction to Probability [M]. Swarthmore Colleg & Dartmouth College, Washington, D. C.: American Mathmetical Society, 1997.
 - [43] Thomas M. Cover, Joy A. Thomas, 信息论基础 [M], 阮吉寿, 张华, 译. 北京: 机械工业出版社, 2008.
 - [44] William Bialek, Ilya Nemenman, and Naftali Tishby. Predictability, Complexity, and Learning[J], Neural Comput. Vol. 13, No. 11: 2409-2463., 2001.
 - [45] Visual Learning with Weak Supersivion, <http://cvn.ecp.fr/tutorials/cvpr2013/>, Organized by Matthew Blaschko, M. Pawan Kumar and Ben Taskar.
 - [46] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg & Demis Hassabis. Human-level control through deep reinforcement learning [J]. Nature vol. 518: 529–533, 2015.
 - [47] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. "Neural Networks for Machine Learning" Lecture 6: Overview of mini-batch gradient descent [OL]. <http://www.cs.toronto.edu/~tijmen/csc321/>

- [48] Marek Rei. 26 Things I Learned In The Deep Learning Summer School [OL]. <http://www.marekrei.com/blog/26-things-i-learned-in-the-deep-learning-summer-school>
- [49] Difference between stochastic gradient descent and online learning [OL]. <https://www.quora.com/Difference-between-stochastic-gradient-descent-and-online-learning>
- [50] Ian J Goodfellow, Oriol Vinyals, and Andrew M Saxe. Qualitatively characterizing neural network optimization problems [OL]. arXiv preprint arXiv:1412.6544, 2014. <https://arxiv.org/abs/1412.6544>
- [51] Choromanska, A., Henaff, M., Mathieu, M., Ben Arous, G., and LeCun, Y. The Loss Surface of Multilayer Networks [OL]. arXiv preprint arXiv:1412.0233, 2014. . <https://arxiv.org/abs/1412.0233>
- [52] David H. Wolpert. Ubiquity symposium: Evolutionary computation and the processes of life: what the no free lunch theorems really mean: how to improve search algorithms [J]. Ubiquity, Volume 2013 Issue December: 1-15, 2013.
- [53] Prateek Gupta. Deep Learning – Regularization [OL]. June 2015. NVIDIA. <https://cs.nyu.edu/mishra/COURSES/15.Summer/L4DNN.pdf>
- [54] Shen Li, Zhouchen Lin, Qingmin Huang, Bastian Leibe, Jiri Matas, Nicu Sebe, Max Welling. Relay Backpropagation for Effective Learning of Deep Convolutional Neural Networks [C]. 14th European Conference, Amsterdam, The Netherlands, October 11–14, Proceedings, Part VII, 2016.
- [55] What is the vanishing gradient problem [OL] . <https://www.quora.com/What-is-the-vanishing-gradient-problem>
- [56] Theano – optimization [OL]: <http://deeplearning.net/software/theano/extending/optimization.html>
- [57] Hugo Larochelle. Foundations of Deep Learning [OL]. Lecture. Université de Sherbrooke, http://www.dmi.usherb.ca/~larocheh/index_en.html
- [58] Andrej Karpathy, Deep Learning for Computer Vision [OL], Bay Area Deep Learning School, 2016. <http://iplab.dmi.unict.it/icvss2016>
- [59] Andrew Gibiansky, Convolutional Neural Network [OL], <http://andrew.gibiansky.com/blog/machine-learning/convolutional-neural-networks/>
- [60] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, "Learning representations by back-propagating errors [J]" Nature Vol. 323 9 October 1986.
- [61] Automatic differentiation [OL], Wikipedia. https://en.wikipedia.org/wiki/Automatic_differentiation
- [62] Boyana Norris, Paul Hovland, Sri Hari Krishna Narayanan, Jean Utke, Automatic Differentiation in Computational Science [C]. Advanced Scientific Computing Advisory Committee Meeting, Nov. 10, 2010.
- [63] Kandel ER, Schwartz JH, Jessell TM. Principles of Neural Science [M], 4th ed. New York: McGraw-Hill, 2000.

-
- [64] S. C. Kremer and J. F. Kolen, A Field Guide to Dynamical Recurrent Neural Networks [M], Singapore: Wiley-IEEE Press, 2001.
 - [65] Vincent Dumoulin and Francesco Visin. A guide to convolution arithmetic for deep learning [OL]. arXiv preprint arXiv:1603.07285, 2016. <https://arxiv.org/abs/1603.07285>
 - [66] Felix Grün, Christian Rupprecht, Nassir Navab, Federico Tombari. A Taxonomy and Library for Visualizing Learned Features in Convolutional Neural Networks [OL]. arXiv preprint arXiv: 1606.07757. <https://arxiv.org/abs/1606.07757>
 - [67] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning Deep Features for Scene Recognition using Places Database [C]. Advances in Neural Information Processing Systems, California: Neural Information Processing Systems Foundation, 2014.
 - [68] Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations [C]. Proceedings of the 26th Annual International Conference on Machine Learning (ICML '09). ACM, New York, NY, USA, 609-616.
 - [69] D. Erhan, Y. Bengio, A. Courville, P. Vincent. Visualizing Higher-Layer Features of a Deep Network [C]. ICML 2009 Workshop on Learning Feature Hierarchies, Montréal, Canada.
 - [70] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. arXiv preprint arXiv:1512.03385, 2015. <https://arxiv.org/abs/1512.03385>
 - [71] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. arXiv preprint arXiv:1603.05027, 2016. <https://www.arxiv.org/abs/1603.05027>
 - [72] Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson. Understanding neural networks through deep visualization [OL]. <http://yosinski.com/deepvis>
 - [73] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks [C]. Proceedings of the European Conference on Computer Vision (ECCV), 2014.
 - [74] Huang, Gao, Liu, Zhuang, and Weinberger, Kilian Q. Densely Connected Convolutional Networks [C]. arXiv:1608.06993, <https://arxiv.org/abs/1608.06993v3>
 - [75] Srivastava, Rupesh K, Greff, Klaus, and Schmidhuber, Juergen. Training Very Deep Networks [J]. In Cortes, C., Lawrence, N. D., Lee, D. D., Sugiyama, M., and Garnett, R., Advances in Neural Information Processing Systems 28, pp. 2377–2385. Curran Associates, Inc., 2015.
 - [76] Srivastava, Rupesh Kumar, Greff, Klaus, and Schmidhuber, Jürgen. Highway Networks [OL]. arXiv:1505.00387, <https://arxiv.org/abs/1505.00387>
 - [77] Veit, Andreas, Wilber, Michael, and Belongie, Serge. Residual Networks are Exponential Ensembles of Relatively Shallow Networks. arXiv:1605.06431, May 2016. <https://arxiv.org/abs/1605.06431>
 - [78] William L. Briggs, Van Emden Henson. A Multigrid Tutorial [OL]. Center for Applied

- Scientific Computing, Lawrence Livermore National Laboratory.
<https://www.math.ust.hk/~mawang/teaching/math532/mgtut.pdf>
- [79] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift [C]. Proceedings of The 32nd International Conference on Machine Learning, pages 448–456, 2015.
- [80] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification [OL]. arXiv preprint arXiv:1502.01852, 2015. <https://arxiv.org/abs/1502.01852>
- [81] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu. Deeply supervised nets [OL]. arXiv preprint arXiv:1409.5185, 2014. <https://arxiv.org/abs/1409.5185>
- [82] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation [C]. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3431–3440, 2015.
- [83] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks [OL]. arXiv preprint arXiv:1211.5063, 2012. <https://arxiv.org/abs/1211.5063>
- [84] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge [OL]. 2014. <http://www.image-net.org/challenges/LSVRC/2014/index>
- [85] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition [OL]. arXiv preprint arXiv:1409.1556, 2014. <https://arxiv.org/abs/1409.1556>
- [86] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions [C]. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1–9, 2015.
- [87] M. Lin, Q. Chen, and S. Yan. Network in network [OL]. arXiv preprint arXiv:1312.4400, 2013. <https://arxiv.org/abs/1312.4400>
- [88] Gustav Larsson, Michael Maire, Gregory Shakhnarovich. FractalNet: Ultra-Deep Neural Networks without Residuals [OL]. arXiv preprint arXiv:1605.07648, May 2016. <https://arxiv.org/abs/1605.07648>
- [89] Sasha Targ, Diogo Almeida, Kevin Lyman, Resnet in Resnet: Generalizing Residual Architectures [OL]. arXiv preprint arXiv: 1603.08029, <http://arxiv.org/abs/1603.08029>.
- [90] L Zhao, J Wang, X Li, Z Tu, W Zeng, On the Connection of Deep Fusion to Ensembling [OL]. arXiv preprint arXiv:1611.07718. <https://arxiv.org/abs/1611.07718>
- [91] Python 2.7 文档 [OL]. <https://docs.python.org/2.7/>
- [92] Hunter, John D. Matplotlib: A 2D graphics environment [J]. Computing in science and engineering Volume: 9, Issue: 3: 90-95. 2007.
- [93] Scipy 文档 [OL]. <https://docs.scipy.org/doc/>
- [94] OpenCV 2.4 文档 [OL]. <http://docs.opencv.org/2.4.13/>
- [95] OpenCV 3.1 文档 [OL]. <http://docs.opencv.org/3.1.0/>

-
- [96] 李沐, MXNet 设计和实现简介 [OL]. <https://github.com/dmlc/mxnet/issues/797>
 - [97] MNIST [OL]. <http://yann.lecun.com/exdb/mnist/>
 - [98] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. FeiFei. Imagenet: A large-scale hierarchical image database [C]. In Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, pages 248–255. IEEE, 2009.
 - [99] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge [J]. International journal of computer vision, 88(2):303–338, 2010.
 - [100] R. B. Girshick. Fast R-CNN [OL]. arXiv preprints, abs/1504.08083, 2015. <https://arxiv.org/abs/1504.08083>
 - [101] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, and S. E. Reed. SSD: single shot multibox detector [OL], abs/1512.02325, 2015. <https://arxiv.org/abs/1512.02325>
 - [102] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection [C]. Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR ‘16): 779-788. IEEE Computer Society, Washington, DC, USA
 - [103] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks [OL]. arXiv preprint arXiv:1506.01497, 2015. <https://arxiv.org/abs/1506.01497>
 - [104] Dan C. Cirean, Alessandro Giusti, Luca M. Gambardella, Jürgen Schmidhuber, "Mitosis Detection in Breast Cancer Histology Images with Deep Neural Networks [J]," Medical Image Computing and Computer-Assisted Intervention – MICCAI 2013 Volume 8150 of the series Lecture Notes in Computer Science pp 411-418.
 - [105] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation [C]. Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR '14). IEEE Computer Society, Washington, DC, USA, 580-587.
 - [106] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets [C]. The British Machine Vision Conference, 2014.
 - [107] J. Uijlings, K. van de Sande, T. Gevers, and A. Smeulders. Selective search for object recognition [J]. International Journal of Computer Vision, Volume 104 (2), page 154-171, 2013
 - [108] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition [J]. IEEE Transactions on Pattern Analysis and Machine Intelligence (Volume: 37, Issue: 9, Sept. 1 2015)
 - [109] J. H. Hosang, R. Benenson, P. Dollar, and B. Schiele. What makes for effective detection proposals? [OL]. arXiv preprint arXiv:1502.05082, 2015.

- <https://arxiv.org/abs/1502.05082>
- [110] Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel. Backpropagation applied to handwritten zip code recognition [J]. *Neural Computation*, Volume 1, Issue 4, Winter 1989, p.541-551.
- [111] S. Chopra, R. Hadsell and Y. LeCun, Learning a similarity metric discriminatively, with application to face verification [C]. 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), 2005, pp. 539-546 vol. 1.
- [112] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. Signature verification using a "Siamese" time delay neural network [C]. In *Proceedings of the 6th International Conference on Neural Information Processing Systems (NIPS'93)*, J. D. Cowan, G. Tesauro, and J. Alspector. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 737-744.
- [113] Leon A. Gatys, Alexander S. Ecker, Matthias Bethge. A Neural Algorithm of Artistic Style [OL]. arXiv:1508.06576. <https://arxiv.org/abs/1508.06576>
- [114] L. A. Gatys, A. S. Ecker and M. Bethge, Image Style Transfer Using Convolutional Neural Networks [C]. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, 2016, pp. 2414-2423.
- [115] L. A. Gatys, A. S. Ecker, and M. Bethge. Texture Synthesis Using Convolutional Neural Networks [C]. In *Advances in Neural Information Processing Systems 28*, 2015.
- [116] A. Mahendran and A. Vedaldi. Understanding Deep Image Representations by Inverting Them [OL]. arXiv preprints. arXiv: 1412.0035. 2014, <https://www.arxiv.org/abs/1412.0035>
- [117] Chuan Li, Michael Wand. Combining Markov Random Fields and Convolutional Neural Networks for Image Synthesis [OL]. arXiv: 1601.04589. <https://arxiv.org/abs/1601.04589>
- [118] Justin Johnson, Alexandre Alahi, Li Fei-Fei. Perceptual Losses for Real-Time Style Transfer and Super-Resolution [OL]. arXiv:1603.08155. <https://arxiv.org/abs/1603.08155>
- [119] Bhautik Joshi, Kristen Stewart, David Shapiro. Bringing Impressionism to Life with Neural Style Transfer in Come Swim [OL]. arXiv:1701.04928. <https://arxiv.org/abs/1701.04928>