CSE 672 Bayesian Vision

University at Buffalo SUNY

Syllabus for Spring 2008

Last updated: 13 Jan 2008

Instructor: Jason Corso (jcorso@cse)

Course Webpage: http://www.cse.buffalo.edu/~jcorso/t/2008spring_vbi.

Syllabus: http://www.cse.buffalo.edu/~jcorso/t/2008spring_vbi/syllabus.pdf.

Downloadable course material can be found on the UBLearns site.

Meeting Times: MW 3-4:20

Location: Bell 242

Office Hours: Tuesday 1-2:30 or by appointment

Main Course Material

Course Overview: The course takes an in-depth look at various Bayesian methods in computer and medical vision. Through the language of Bayesian inference, the course will present a coherent view of the approaches to various key problems such as detecting objects in images, segmenting object boundaries, and recognizing objects. The course is roughly partitioned into two parts: modeling and inference. In the first half, it will cover both classical models such as weak membrane models and Markov random fields as well as more recent models such as conditional random fields, latent Dirichlet allocation, and topic models. In the second half, it will focus on inference algorithms. Methods include PDE boundary evolution algorithms such as region competition, discrete optimization methods such as graph-cuts and graph-shifts, and stochastic optimization methods such as data-driven Markov chain Monte Carlo. An emphasis will be placed on both the theoretical aspects of this field as well as the practical application of the models and inference algorithms.

Course Project: Each student will be required to implement a course project that is either a direct implementation of a method discussed during the semester or new research in Bayesian vision. A paper describing the project is required at the end of the semester (8-10 pages two column IEEE format) and we will have an open-house poster session to present the projects. Working project demos are suggested but not required for the poster session.

Prerequisites: It is assumed that the students have taken introductory courses in pattern recognition (CSE 555), and computer vision (CSE 573). Machine learning (CSE 574) is suggested but not required. Permission of the instructor is required if these pre-requisites have not been met.

Course Goals: After taking the course, the student should have a clear understanding of the state-of-the-art models and inference algorithms for solving vision problems within a Bayesian methodology. Through completing the course project, the student will also have a deep understanding of the low-level details of a particular model/algorithm.

Textbooks: There is unfortunately no complete textbook for this course. The required material will either be distributed by the instructor or found on reserve at the UB Library. Recommended textbooks are

1. Li, S. *Markov Random Field Modeling in Image Analysis*. Springer-Verlag. 2001.
2. Winkler, G. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods: A Mathematical Introduction*. Springer. 2006.

3. Chalmond, B. *Modeling and Inverse Problems in Image Analysis*. Springer. 2003.
4. Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer. 2007.

Grading: Letter grading distributed as follows:

- Discussion (20%)
- Homeworks (30%)
- Project (50%)

Homeworks: There will be three homeworks, equally weighted. They will cover both theoretical and practical (implementation) aspects of the material. Students may collectively discuss the homework problems, but they must write them independently. No sharing of written/typed materials of any sort is allowed.

Programming Language: Student choice for homeworks and project (generally, Matlab, Java, or C/C++). However, no platform-specific libraries/packages are permissible.

Course Outline

The course is roughly divided into two parts. In the first part, we discuss various modeling and associated learning algorithms. In the second part, we discuss the computing and inference algorithms which use the previously discussed models to solve complex inference problems in vision. The topic outline follows; citations are given and an underlined citation indicates a primary (must-read) one.

1. Introduction.
   (a) Discussion of Bayesian inference in the context of vision problems. (Winkler, 2006, Chapter 1) (Chalmond, 2003, Chapter 1) (Hanson, 1993)
   (b) Presentation of relevant empirical findings concerning the statistics of images motivating the Bayesian approach. (Field, 1994) (Field, 1987) (Julesz, 1981) (Kersten, 1987) (Ruderman, 1994) (Simoncelli & Olshausen, 2001) (Torralba & Oliva, 2003) (Wu *et al.*, 2007)
   (c) Model classes: discriminative, generative and descriptive. (Zhu, 2003)
2. Modeling and Learning.
   (a) Descriptive models on regular lattices.
      i. Markov random field models and Gibbs fields. (Li, 2001, §1.2) (Winkler, 2006, §2,3) (Dubes & Jain, 1989)
      ii. The Hammersley-Clifford theorem.
      iii. Bayes MRF Estimators (Winkler, 2006, §1.4) (Li, 2001, §1.5) (Geman & Geman, 1984)
      iv. Examples:
         A. Auto-Models (Besag, 1974) (Li, 2001, §1.3.1, 2.3, 2.4) (Winkler, 2006, §15)
         B. Weak membrane models, Mumford-Shah, TV, etc.
      v. Applications:
         A. Image Restoration and Denoising (Li, 2001, §2.2)
         B. Edge Detection and Line Processes (Li, 2001, §2.3) (Geman & Geman, 1984)
         C. Texture (Li, 2001, §2.4) (Winkler, 2006, §15,16)
      vi. MRF Parameter Estimation (Li, 2001, §6) (Winkler, 2006, §5,6)
         A. Maximum-Likelihood
         B. Pseudo-Likelihood
         C. Gibbs Sampler (and brief introduction to MCMC)
   (b) Descriptive Models on Regular Lattices: Advanced Topics
      i. Discontinuities and Smoothness Priors (Li, 2001, §4)
      ii. FRAME and Minimax entropy learning of potential functionals. (Zhu *et al.*, 1998) (Zhu *et al.*, 1997) (Coughlan & Yuille, 2003)
      iii. Hidden Markov random fields. (Zhang *et al.*, 2001)

iv. Conditional random fields. (Lafferty *et al.*, 2001) (Kumar & Hebert, 2003) (Wallach, 2004)
v. MRF as a foundation for multiresolution computing. (Gidas, 1989)

(c) Descriptive and Generative Models on Irregular Graphs and Hierarchies.

i. Markov random field hierarchies. (Derin & Elliott, 1987) (Krishnamachari & Chellappa, 1995) (Chardin & Perez, 1999)
ii. Over-Complete Bases and Sparse Coding (Zhu, 2003, §6) (Olshausen & Field, 1997) (Coifman & Wickerhauser, 1992)
iii. Textons (Julesz, 1981) (Zhu *et al.*, 2005) (Malik *et al.*, 1999)
iv. And-Or graphs and context-sensitive grammars. (Zhu & Mumford, 2006) (Han & Zhu, 2005)
v. Dirichlet Processes (DP) and Bayesian Clustering (Ferguson, 1973)
vi. Latent Dirichlet Allocation, hierarchical DP and author-topic models. (Blei *et al.*, 3) (Teh *et al.*, 2005) (Steyvers *et al.*, 2004)
vii. Correspondence LDA (Blei & Jordan, 2003)

(d) Integrating Descriptive and Generative Models (Guo *et al.*, 2006)

3. Inference Algorithms.

(a) Boundary methods.

i. Level set evolution. (Chan & Vese, 2001)
ii. Region competition algorithm. (Zhu & Yuille, 1996a)

(b) Discrete Deterministic Inference.

i. Graph-Cuts: α-Expansion algorithm and min-cut/max-flow relationship. (Boykov *et al.*, 2001) (Kolmogorov & Zabih, 2002a)
ii. Graph-Shifts algorithm. (Corso *et al.*, 2007) (Corso *et al.*, 2008)
iii. Sum-Product algorithm (exact Belief Propagation). (Bishop, 2006, §8) (Yedidia *et al.*, 2001) (Frey & MacKay, 1997) (Felzenszwalb & Huttenlocher, 2006)
iv. Generalized Belief Propagation. (Yedidia *et al.*, 2005) (Yedidia *et al.*, 2000)
v. Inference on And-Or graphs. (Zhu & Mumford, 2006) (Han & Zhu, 2005)

(c) Stochastic Inference. (Forsyth *et al.*, 2001)

i. Gibbs sampling. (Geman & Geman, 1984) (Winkler, 2006, §5,7)
ii. Metropolis-Hastings and Markov chain Monte Carlo methods. (Winkler, 2006, §10) (Tierney, 1994) (Liu, 2002)
iii. Data-Driven MarkovMCMC algorithm. (Tu & Zhu, 2002) (Tu *et al.*, 2005) (Green, 1995)
iv. Swendsen-Wang algorithm. (Swendsen & Wang, 1987) (Barbu & Zhu, 2005) (Barbu & Zhu, 2004)
v. Sequential MCMC and Particle Filters. (Isard & Blake, 1998) (Liu & Chen, 1998)

Calendar

The calendar will be populated as the semester proceeds based on the above course outline and our progress.

	Week	Monday	Wednesday	Assignments
1	1/14	Introduction. Vision as Bayesian Inference.	Statistics of Natural Images	
2	1/21	**MLK Day – No Class**	Introduction to MRF Models and Gibbs Fields.	
3	1/28			
4	2/4			
5	2/11			
6	2/18			
7	2/25			
8	3/3	Author-Topic Models and Latent-Dirichlet Allocation.		
	3/10	**Spring Recess – No Classes**		

9	3/17			
10	3/24			
11	3/31			
12	4/7			
13	4/14			
14	4/21			
15	4/28		Public Poster / Demo Session for Projects	

Project

The goal of the project is to have each student (or pair of students) study and implement one particular model-inference pair. Below is a list of possible projects, but the student is encouraged to design a project of their own device. The ultimate goal is for each student to do some new work. Within reason, camera and video equipment will be made available to the students from the Vision and Perceptual Machines Lab (Dr. Corso's Lab in the Commons CSE Research Area). Suitable arrangements should be made with the instructor to facilitate equipment use.

List of Possible Projects

- Learning and sampling generic image priors such a line processes (1).
- MRF Potential Learning by Minimax Entropy (1).
- Sampling Julesz ensemble of textures (1).
- Inference by Tree-Reweighted Message Passing (1).
- Learning and sampling a stochastic graph model (2).
- Learning and sampling the primal sketch from natural or medical images (2).

Project Schedule

2/4 Project proposal due in class. 1-page description of the proposed project and the type of problem/data. It should include three milestones in planning.

3/5 Milestone 1 Report due in class. (1-paragraph)

3/31 Milestone 2 Report due in class. (1-paragraph)

4/30 Final milestone and public poster / demo session.

5/5 9AM Project write-up and source code are due.

Project Write-Up

The write-up will be in standard two-column IEEE journal format at a maximum of 10 pages. It should be approached as a standard paper containing introduction and related work, methodology, results, and discussion.

Additional Information

Similar Courses at Other Institutions: (incomplete and in no important order)

- Professor Alan Yuille at UCLA. *Vision as Bayesian Inference*. http://www.stat.ucla.edu/~yuille/courses/Stat238/Stat_238.htm
- Professor Song-Chun Zhu at UCLA. *Statistical Modeling and Learning in Vision and Image Science*. http://www.stat.ucla.edu/%7Esczhu/Courses/UCLA/Stat_232A/Stat_232A.html
- Professor Song-Chun Zhu at UCLA. *Statistical Computing and Inference in Vision and Image Science*. http://www.stat.ucla.edu/%7Esczhu/Courses/UCLA/Stat_232B/Stat_232B.html
- Professor Fei-Fei Li at Princeton. *High-Level Recognition in Computer Vision* http://vision.cs.princeton.edu/cs598_spring07/
- Professor Tal Arbel at McGill. *Statistical Computer Vision* http://www.cim.mcgill.ca/~arbel/courses/626.html
- Professor William T. Freeman at MIT. *Advances in Computer Vision: Learning and Interfaces* http://courses.csail.mit.edu/6.869/

Course Bibliography

Most items below have been cited above, but there are also some additional references that extend the content of the course. When available, PDFs of articles have been uploaded to the UBLearns "Course Documents" section. The naming convention is the first two characters of (up to) the first three authors following by an acronym for the venue (e.g., CVPR for Computer Vision and Pattern Recognition) followed by the year. So, the Geman and Geman 1984 PAMI article is GeGePAMI1984.pdf.

References

Barbu, A., & Zhu, S. C. 2004. Multigrid and Multi-level Swendsen-Wang Cuts for Hierarchic Graph Partitions. *Pages 731–738 of: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2. 3

Barbu, A., & Zhu, S. C. 2005. Generalizing Swendsen-Wang to Sampling Arbitrary Posterior Probabilities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **27**(8), 1239–1253. 3

Besag, J. 1974. Spatial interaction and the statistical analysis of lattice systems (with discussion). *J. Royal Stat. Soc., B*, **36**, 192–236. 2

Besag, J. 1986. On The Statistical Analysis of Dirty Pictures (with discussion). *Journal of the Royal Statistical Society [Ser. B]*, **48**, 259–302.

Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*. Springer. 3

Bishop, C. M., & Winn, J. M. 2000. Non-linear Bayesian Image Modelling. *Pages 3–17 of: European Conference on Computer Vision*, vol. 1.

Blei, D. M., & Jordan, M. I. 2003. Modeling Annotated Data. *In: Proceedings of SIGIR*. 3

Blei, D. M., Ng, A. Y., & Jordan, M. I. 3. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, **2003**, 993–1022. 3

Bouman, C. A., & Shapiro, M. 1994. A multiscale random field model for Bayesian image segmentation. *Image Processing, IEEE Transactions on*, **3**(2), 162–177.

Boykov, Y., Veksler, O., & Zabih, R. 2001. Fast Approximate Energy Minimization via Graph Cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **23**(11), 1222–1239. 3

Chalmond, B. 2003. *Modeling and Inverse Problems in Image Analysis*. Applied Mathematical Sciences, vol. 155. Springer. 2

Chan, T. F., & Vese, L. A. 2001. Active Contours Without Edges. *IEEE Transactions on Image Processing*, **10**(2), 266–277. 3

Chardin, A., & Perez, P. 1999. Semi-iterative Inferences with Hierarchical Energy-Based Models for Image Analysis. *Energy Minimization Methods in Computer Vision and Pattern Recognition: Second International Workshop, EMMCVPR'99, York, UK, July 1999. Proceedings*, 730–730. 3

Coifman, R. R., & Wickerhauser, M. V. 1992. Entropy-based Algorithms for Best Basis Selection. *IEEE Transactions on Information Theory*, **38**(2), 713–718. 3

Cootes, T.F., & Taylor, C.J. 2004. *Statistical Models of Appearance for Computer Vision*. Tech. rept. Imaging Science and Biomedical Engineering, University of Manchester.

Corso, J. J., Sharon, E., Dube, S., El-Saden, S., Sinha, U., & Yuille, A. Efficient Multilevel Brain Tumor Segmentation with Integrated Bayesian Model Classification. *IEEE Transactions on Medical Imaging*. (in press).

Corso, J. J., Sharon, E., & Yuille, A. 2006. Multilevel Segmentation and Integrated Bayesian Model Classification with an Application to Brain Tumor Segmentation. *Pages 790–798 of: Medical Image Computing and Computer Assisted Intervention*, vol. 2.

Corso, J. J., Tu, Z., Yuille, A., & Toga, A. W. 2007. Segmentation of Sub-Cortical Structures by the Graph-Shifts Algorithm. *Pages 183–197 of:* Karssemeijer, N., & Lelieveldt, B. (eds), *Proceedings of Information Processing in Medical Imaging*. 3

Corso, J. J., Tu, Z., & Yuille, A. 2008. MRF Labeling with a Graph-Shifts Algorithm. *In: Proceedings of International Workshop on Combinatorial Image Analysis*. (to appear). 3

Coughlan, J. M., & Yuille, A. L. 2003. Algorithms from Statistical Physics for Generative Models of Images. *Image and Vision Computing, Special Issue on Generative-Model Based Vision*, **21**(1), 29–36. 2

Derin, H., & Elliott, H. 1987. Modeling and Segmentation of Noisy and Texture Images Using Gibbs Random Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **9**(1), 39–55. 3

Dubes, R. C., & Jain, A. K. 1989. Random field models in image analysis. *Journal of Applied Statistics*, **16**(2), 131 – 164. 2

Fei-Fei, L., & Perona, P. 2005. A Bayesian Hierarchical Model for Learning Natural Scene Categories. *In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*.

Felzenszwalb, P. F., & Huttenlocher, D. P. 2006. Efficient Belief Propagation for Early Vision. *International Journal of Computer Vision*, **70**(1). 3

Ferguson, T. S. 1973. A Bayesian Analysis of Some Nonparametric Problems. *The Annals of Statistics*, **1**(2), 209–230. 3

Field, D. J. 1987. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, **4**(12), 2379–2394. 2

Field, D. J. 1994. What is the goal of sensory coding? *Neural Computation*, **6**, 559–601. 2

Forsyth, D., Haddon, J., & Ioffe, S. 2001. The Joy of Sampling. *International Journal of Computer Vision*, **41**(1), 109–134. 3

Frey, B. J., & MacKay, D. 1997. A Revolution: Belief Propagation in Graphs with Cycles. *In: Proceedings of Neural Information Processing Systems (NIPS)*. 3

Geman, S., & Geman, D. 1984. Stochastic Relaxation, Gibbs Distributions, and Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**, 721–741. 2, 3

Gidas, B. 1989. A Renormalization Group Approach to Image Processing Problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **11**(2), 164–180. 3

Green, P. J. 1995. Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination. *Biometrika*, **82**(4), 711–732. 3

Guo, C. E., Zhu, S. C., & Wu, Y. N. 2003. Modeling Visual Patterns by Integrating Descriptive and Generative Models. *International Journal of Computer Vision*, **53**(1), 5–29.

Guo, C. E., Zhu, S. C., & Wu, Y. N. 2006. Primal Sketch: Integrating Texture and Structure. *Computer Vision and Image Understanding*. (to appear). 3

Han, F., & Zhu, S. C. 2005. Bottom-Up/Top-Down Image Parsing by Attribute Graph Grammar. *Pages 1778–1785 of: Proceedings of International Conference on Computer Vision*, vol. 2. 3

Hanson, K. M. 1993. Introduction to Bayesian image analysis. *Medical Imaging: Image Processing*, **Proc. SPIE 1898**, 716–731. 2

Held, K., Kops, E. R., Krause, B. J., Wells, W. M., III., Kikinis, R., & Muller-Gartner, H. W. 1997. Markov random field segmentation of brain MR images. *Medical Imaging, IEEE Transactions on*, **16**(6), 878–886.

Isard, M., & Blake, A. 1998. CONDENSATION – conditional density propagation for visual tracking. *International Journal of Computer Vision*, **29**(1), 5–28. 3

Julesz, B. 1981. Textons, the elements of texture perception and their interactions. *Nature*, **290**, 91–97. 2, 3

Kersten, D. 1987. Predictability and Redundancy of Natural Images. *Journal of the Optical Society of America*, **A 4**(12), 2395–2400. 2

Kolmogorov, V., & Zabih, R. 2002a. What Energy Functions Can Be Minimized via Graph Cuts? *Pages 65–81 of: European Conference on Computer Vision*, vol. 3. 3

Kolmogorov, Vladimir, & Zabih, Ramin. 2002b. Multicamera Scene Reconstruction via Graph-Cuts. *Pages 82–96 of: European Conference on Computer Vision*.

Krishnamachari, S., & Chellappa, R. 1995. Multiresolution GMRF models for texture segmentation. vol. 4. 3

Kumar, S., & Hebert, M. 2003. Discriminative Random Fields: A Discriminative Framework for Contextual Interaction in Classification. *In: International Conference on Computer Vision*. 3

Lafferty, J., McCallum, A., & Pereira, F. 2001. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. *In: Proceedings of International Conference on Machine Learning*. 3

Lee, C. H., Schmidt, M., Murtha, A., Bistritz, A., Sander, J., & Greiner, R. 2005. Segmenting Brain Tumor with Conditional Random Fields and Support Vector Machines. *In: Proceedings of Workshop on Computer Vision for Biomedical Image Applications at International Conference on Computer Vision*.

Lee, S., & Crawford, M. M. 2005. Unsupervised multistage image classification using hierarchical clustering with a bayesian similarity measure. *Image Processing, IEEE Transactions on*, **14**(3), 312–320.

Li, S. Z. 2001. *Markov Random Field Modeling in Image Analysis*. 2nd edn. Springer-Verlag. 2

Liu, J. S. 2002. *Monte Carlo Strategies in Scientific Computing*. Springer. 3

Liu, J. S., & Chen, R. 1998. Sequential Monte Carlo Methods for Dynamic Systems. *Journal of the American Statistical Society*, **93**(443), 1032–1044. 3

MacEachern, S. N., & Muller, P. 1998. Estimating Mixture of Dirichlet Process Models. *Journal of Computational and Graphical Statistics*, **7**(2), 223–238.

Malik, Jitendra, Belongie, Serge, Shi, Jianbo, & Leung, Thomas. 1999. Textons, Contours, and Regions: Cue Combination in Image Segmentation. *In: International Conference on Computer Vision*. 3

Naphade, M. R., & Huang, T. S. 2001. A Probabilistic Framework for Semantic Video Indexing, Filtering, and Retrieval. *IEEE Transactions on Multimedia*, **3**(1), 141–151.

Olshausen, B. A., & Field, D. J. 1997. Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1? *Vision Research*, **37**(23), 3311–3325. 3

Raj, A., & Zabih, R. 2005. A Graph Cut Algorithm for Generalized Image Deconvolution. *In: Proceedings of International Conference on Computer Vision*.

Ranganathan, A. 2004 (September). *The Dirichlet Process Mixture (DPM) Model*.

Richardson, S., & Green, P. J. 1997. On Bayesian Analysis of Mixtures With an Unknown Number of Components. *Journal of the Royal Statistical Society – Series B*, **59**(4), 731–758.

Ruderman, D. L. 1994. The statistics of natural images. *Network: Computation in Neural Systems*, **5**(4), 517–548. 2

Schaap, M., Smal, I., Metz, C., Walsum, T. van, & Niessen, W. 2007. Bayesian Tracking of Elongated Structures in 3D Images. *In:* Karssemeijer, N., & Lelieveldt, B. (eds), *Proceedings of Information Processing in Medical Imaging*.

Simoncelli, E. P., & Olshausen, B. A. 2001. Natural Image Statistics and Neural Representation. *Annual Review of Neuroscience*, **24**, 1193–1216. 2

Steyvers, M., Smyth, P., Rosen-Zvi, M., & Griffiths, T. 2004. Probabilistic Author-Topic Models for Information Discovery. *In: 10th ACM SigKDD Conference on Knowledge Discovery and Data Mining*. 3

Sudderth, E. B., Torralba, A., Freeman, W. T., & Willsky, A. S. 2005. Describing Visual Scenes using Transformed Dirichlet Processes. *In: Proceedings of Neural Information Processing Systems (NIPS)*.

Swendsen, R. H., & Wang, J. S. 1987. Nonuniversal Critical Dynamics in Monte Carlo Simulations. *Physical Review Letters*, **58**(2), 86–88. 3

Teh, Y. W., Jordan, M. I., Beal, M. J., & Blei, D. M. 2005. Hierarchical Dirichlet Processes. *In: Advances in Neural Information Processing Systems (NIPS) 17*. 3

Tierney, L. 1994. Markov Chains for Exploring Posterior Distributions. *The Annals of Statistics*, **22**(4), 1701–1728. 3

Torr, Phil, & Davidson, C. 2003. IMPSAC: Synthesis of Importance Sampling and Random Sample Consensus. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **25**(3), 354–364.

Torralba, Antonio, & Oliva, A. 2003. Statistics of Natural Image Categories. *Network: Computation in Neural Systems*, **14**, 391–412. 2

Torre, F., & Black, M. J. 2001. Robust Principal Component Analysis for Computer Vision. *In: International Conference on Computer Vision*.

Tu, Z., & Zhu, S. C. 2002. Image Segmentation by Data-Driven Markov Chain Monte Carlo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **24**(5), 657–673. 3

Tu, Z., Chen, X. R., Yuille, A. L., & Zhu, S. C. 2005. Image Parsing: Unifying Segmentation, Detection and Recognition. *International Journal of Computer Vision*. 3

Wallach, H. M. 2004. *Conditional Random Fields: An Introduction*. CIS MS-CIS-04-21. University of Pennsylvania. 3

Winkler, G. 2006. *Image Analysis, Random Fields, and Markov Chain Monte Carlo Methods*. 2nd edn. Springer. 2, 3

Wu, Y. N., Zhu, S. C., & Guo, C. E. 2007. From Information Scaling of Natural Images to Regimes of Statistical Models. *Quarterly of Applied Mathematics*. 2

Yedidia, J. S., Freeman, W. T., & Weiss, Y. 2000. Generalized Belief Propagation. *Pages 689–695 of: Advances in Neural Information Processing Systems (NIPS)*, vol. 13. 3

Yedidia, J. S., Freeman, W. T., & Weiss, Y. 2001 (May). *Bethe free energy, Kikuchi approximations and belief propagation algorithms*. Tech. rept. TR2001-16. Mitsubishi Electronic Research Laboratories. 3

Yedidia, J. S., Freeman, W. T., & Weiss, Y. 2005. Constructing Free-Energy Approximations and Generalized Belief Propagation Algorithms. *IEEE Transactions on Information Theory*, **51**(7), 2282–2312. 3

Zabih, Ramin, & Kolmogorov, Vladimir. 2004. Spatially Coherent Clustering Using Graph Cuts. *Pages 437–444 of: IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2.

Zhang, Y., Brady, M., & Smith, S. 2001. Segmentation of Brain MR Images Through a Hidden Markov Random Field Model and the Expectation-Maximization Algorithm. *IEEE Transactions on Medical Imaging*, **20**(1), 45–57. 2

Zhu, S. C. 1999. Stochastic Jump-Diffusion Process for Computing Medial Axes in Markov Random Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**(11), 1158–1169.

Zhu, S. C. 2003. Statistical Modeling and Conceptualization of Visual Patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **25**(6), 691–712. 2, 3

Zhu, S. C., & Mumford, D. 1997. Prior Learning and Gibbs Reaction-Diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **19**(11), 1236–1250.

Zhu, S. C., & Mumford, D. 2006. A Stochastic Grammar of Images. *Foundations and Trends in Computer Graphics and Vision*, **2**(4), 259–362. 3

Zhu, S. C., & Yuille, A. 1996a. Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for Multiband Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**(9), 884–900. 3

Zhu, S. C., & Yuille, Alan L. 1996b. FORMS: A Flexible Object Recognition and Modeling System. *International Journal of Computer Vision*, **20**(3), 187–212.

Zhu, S. C., Wu, Y., & Mumford, D. 1997. Minimax Entropy Principle and Its Application to Texture Modeling. *Neural Computation*, **9**(8), 1627–1660. 2

Zhu, S. C., Wu, Y. N., & Mumford, D. B. 1998. FRAME: Filters, Random field And Maximum Entropy: — Towards a Unified Theory for Texture Modeling. *International Journal of Computer Vision*, **27**(2), 1–20. 2

Zhu, S. C., Guo, C. E., Wang, Y., & Xu, Z. 2005. What are Textons? *International Journal of Computer Vision*, **62**(1), 121–143. 3