CSE 672 Bayesian Vision

SUNY at Buffalo

Syllabus for Fall 2012

Last updated: 24 Aug 2012

Instructor: Jason Corso (UBIT: jcorso)

Course Webpage: `http://www.cse.buffalo.edu/~jcorso/t/CSE672`.

Syllabus: `http://www.cse.buffalo.edu/~jcorso/t/CSE672/files/syllabus.pdf`.

Downloadable course material can be found on the CSE UNIX network: `/home/csefaculty/jcorso/672`.

Meeting Times: TR 12:30-1:50

Location: Fronczak Hall 422 (`http://goo.gl/maps/0Y4mz`)

Email Listserv: `cse672-fa12-list@listserv.buffalo.edu`
Use this list for any and all course discussion, except private matters.

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 activities in video. 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, 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 near the end of the semester (6-8 pages two column IEEE format). The papers will be peer-reviewed in the course; revisions need to be made based on the peer review and the final submission needs to include a *letter to the editor* describing the paper as if it is in submission to a journal and a description of the revisions made and why. Working project demos are required at the end of the semester. **This is a "projects" course. Your projects can satisfy a Masters requirement. In most cases, it will involve at least some new/independent research. Previous offerings of this course have resulted in numerous papers accepted at major conferences and journals.**

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. A strong understanding and ability to work with probabilities, statistics, calculus and optimization is expected.

Permission of the instructor is required if these pre-requisites have not been met.

Course Goals: After taking the course, the student should will 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 and application. The student will have completed some independent research in Bayesian Vision by the end of the course.

The student will also have experience in planning a project, conducting semi-independent research, and writing up the results; peer-review practice will also be part of the course.

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 below; it is suggested you pick up a copy of at least one of the first three (and if all students do this there will be a half dozen copies of each floating around to share).

1. Li, S. *Markov Random Field Modeling in Image Analysis*. Springer-Verlag. 3rd Edition. 2009.
2. Winkler, G. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods: A Mathematical Introduction*. Springer. 2006.
3. Blake, A., Kohli, P. and Rother, C. *Markov Random Fields for Vision and Image Processing*. MIT Press. 2011.
4. Chalmond, B. *Modeling and Inverse Problems in Image Analysis*. Springer. 2003.
5. Koller, D. and Friedman, N. *Probabilistic Graphical Models: Principles and Techniques.* MIT Press. 2009.
6. Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer. 2007.

Grading: Letter grading distributed as follows:

- In-Class Discussion/Quizzing (50%)
- Homeworks (0%)
- Project (50%)

In-Class Discussion/Quizzing: Half of the grades in this course are based on the students (1) participation in the class, (2) ability to answer questions when queried and (3) ask questions. No written quizzes are planned, but the professor reserves the possibility.

Homeworks: There will be weekly homeworks recommended. They will cover both theoretical and practical (implementation) aspects of the material. The homework assignments are not turned in. We will organize a weekly time where the students in the course will come together to discuss the weekly work without the professor around.

Programming Language: Student choice for the project (generally, Python, Matlab, Java, or C/C++). Any course-relevant aspects of the project need to be independently developed; e.g., if you are using belief propagation as your project's inference algorithm, then you need to implement belief propagation from scratch. No exceptions; don't ask.

For the homeworks and some in-class exercises we will use the UGM library written by Dr. Mark Schmidt; `http://www.di.ens.fr/~mschmidt/Software/UGM.html`. At various points in the course, you will be asked to either run through a demo/function from the library or implement/reimplement a different method for pedagogical value. There will be no introduction to the library in the course, you are expected to learn it in the first week or two (work through the early and simple demos "Small," "Chain," "Tree", and "ICM."

Word Processing: This course forces you to learn LaTeX if you do not already know it. It is the language of the realm. All things submitted to me must be formatted in LaTeX.

Working 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. All or most papers are available in PDF at the course directory (location above).

Paper citations are given below (somewhat sparsely), but few references are given to chapters in the books mentioned above. It is suggested you look in the books for more information when needed.

1. Introduction.
   (a) Discussion of Bayesian inference in the context of vision problems. <u>[Winkler, 2006, Chapter 1]</u> [Chalmond, 2003, Chapter 1] [Hanson, 1993] Probabilistic Inference Primer: [Griffiths and Yuille, 2006]
   (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 and Olshausen, 2001] [Torralba and Oliva, 2003] [Wu et al., 2007]
   (c) Model classes: discriminative, generative and descriptive. <u>[Zhu, 2003]</u>
2. Modeling and Learning.
   (a) Descriptive models on regular lattices.
      i. Markov random field models and Gibbs fields. <u>[Li, 2001, §1.2]</u> <u>[Winkler, 2006, §2,3]</u> [Dubes and Jain, 1989]
      ii. The Hammersley-Clifford theorem.
      iii. Bayes MRF Estimators [Winkler, 2006, §1.4] [Li, 2001, §1.5] [Geman and 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 and 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)
D. Large Margin Methods [Blake et al., 2011, §15]

(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 and Yuille, 2003]
iii. Hidden Markov random fields. [Zhang et al., 2001]
iv. Conditional random fields. [Lafferty et al., 2001] [Kumar and Hebert, 2003] [Wallach, 2004] [Ladicky et al., 2009]
v. MRF as a foundation for multiresolution computing. [Gidas, 1989]
vi. Higher Order Extensions [Kohli et al., 2007] [Kohli et al., 2009] and Field of Experts [Roth and Black, 2009].

(c) Descriptive and Generative Models on Irregular Graphs and Hierarchies.
i. Markov random field hierarchies. [Derin and Elliott, 1987] [Krishnamachari and Chellappa, 1995] [Chardin and Perez, 1999]
ii. Over-Complete Bases and Sparse Coding [Zhu, 2003, §6] [Olshausen and Field, 1997] [Coifman and Wickerhauser, 1992]
iii. Textons [Julesz, 1981] [Zhu et al., 2005] [Malik et al., 1999]
iv. And-Or graphs and context-sensitive grammars. [Zhu and Mumford, 2007] [Han and Zhu, 2005]
v. Dirichlet Processes (DP) and Bayesian Clustering [Ferguson, 1973]
vi. Latent Dirichlet Allocation, hierarchical DP and author-topic models. [Blei et al., 2003] [Teh et al., 2005] [Steyvers et al., 2004]
vii. Correspondence LDA [Blei and Jordan, 2003]

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

3. Inference Algorithms.

(a) Boundary methods.
i. Level set evolution. [Chan and Vese, 2001]
ii. Region competition algorithm. [Zhu and Yuille, 1996a]

(b) Exact Inference. *Exploit the structure of the graph or the form of the potentials to search for the global optimum efficiently (in polynomial time).*
i. Chains and Trees.
ii. Sum-Product algorithm (exact Belief Propagation). [Bishop, 2006, §8] [Yedidia et al., 2001] [Frey and MacKay, 1997] [Felzenszwalb and Huttenlocher, 2006]
iii. Graph-Cuts: min-cut/max-flow relationship. [Blake et al., 2011, §2]
What energy functions can/can not be minimized by graph cuts? [Kolmogorov and Zabih, 2004]

(c) Approximate Inference.
i. Discrete Deterministic Inference.
A. Graph-Cuts: α-Expansion algorithm. [Boykov et al., 2001]
B. Graph-Shifts algorithm. [Corso et al., 2007] [Corso et al., 2008b]
C. Generalized Belief Propagation. [Yedidia et al., 2005] [Yedidia et al., 2000]
D. Inference on And-Or graphs. [Zhu and Mumford, 2007] [Han and Zhu, 2005]

ii. Stochastic Inference. [Forsyth et al., 2001]
A. Mean Field Approximation.
B. Gibbs sampling. [Geman and Geman, 1984] [Winkler, 2006, §5,7]
C. Metropolis-Hastings and Markov chain Monte Carlo methods. [Winkler, 2006, §10] [Tierney, 1994] [Liu, 2002]
D. Data-Driven MarkovMCMC algorithm. [Tu and Zhu, 2002] [Tu et al., 2005] [Green, 1995]
E. Swendsen-Wang algorithm. [Swendsen and Wang, 1987] [Barbu and Zhu, 2005] [Barbu and Zhu, 2004]
F. Sequential MCMC and Particle Filters. [Isard and Blake, 1998] [Liu and Chen, 1998]

Project

The goal of the project is to have each student solve a real problem using the ideas learned herein. The professor will distribute/discuss project ideas in the first week of the class; students are encouraged to design their own project in conjunction with the professor. The ultimate goal is for each student to do some new work and learn by doing so. Within reason, camera and video equipment will be made available to the students from the VPML (my lab). Suitable arrangements should be made with the instructor to facilitate equipment use.

Project topics can cover a myriad of current problems in vision and must include some technical aspect developed on top of ideas in the course. A project focusing on statistics of a class of images/videos is also fair game but will need to be thoroughly justified.

Project Schedule

9/11 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. (All writing must be done in LaTeX.)

9/20 Project plan due in class (this is the refinement of the project proposal; i.e., project proposal v 2). 3-page description of the proposed project, the most related work from the literature, the three milestones, planned data and experiments, and *a goal statement* that presents a table with two columns:

Outcome	Grade
My project will blah blah blah	A
My project will blah blah blah	B
My project will blah blah blah	C
My project will not work.	F

You fill in the blah blah blah and I'll consider it (and approve it or make you modify it). Hence, your Project plan is a contract and you have just graded yourself.

10/4 Milestone 1 Report due in class. (1-paragraph)

10/18 Milestone 2 Paper due in class. (4ish-pages)

11/1 Milestone 3 Paper due in class. (full paper)

11/1-6 Blind Peer Review Period. (Round robin with everyone reviewing two papers.)

11/15 Revised paper due. (Note 11/15 is the CVPR deadline.)

after 11/15 Project presentations and demos in class.

Project Write-Up

The paper should be in standard IEEE conference format at a maximum of 8 pages. We'll explain in class how to set it up. 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.

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

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

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

J. Besag. On the statistical analysis of dirty pictures (with discussion). *Journal of the Royal Statistical Society [Ser. B]*, 48:259–302, 1986.

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

C. M. Bishop and J. M. Winn. Non-linear Bayesian Image Modelling. In *European Conference on Computer Vision*, volume 1, pages 3–17, 2000.

A. Blake, P. Kohli, and C. Rother, editors. *Markov Random Fields for Vision and Image Processing*. MIT Press, 2011. 3

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

D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003. 3

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

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

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

T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. on Image Processing*, 10(2):266–277, 2001. 3

A. Chardin and P. Perez. 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*, pages 730–730, 1999. URL `http://www.springerlink.com/content/6yq1rglku6ccxjpu`. 3

R. R. Coifman and M. V. Wickerhauser. Entropy-based algorithms for best basis selection. *IEEE Transactions on Information Theory*, 38(2):713–718, 1992. 3

T.F. Cootes and C.J. Taylor. Statistical Models of Appearance for Computer Vision. Technical report, Imaging Science and Biomedical Engineering, University of Manchester, 2004.

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

J. J. Corso, Z. Tu, A. Yuille, and A. W. Toga. Segmentation of Sub-Cortical Structures by the Graph-Shifts Algorithm. In N. Karssemeijer and B. Lelieveldt, editors, *Proceedings of Information Processing in Medical Imaging*, pages 183–197, 2007. 3

J. J. Corso, E. Sharon, S. Dube, S. El-Saden, U. Sinha, and A. Yuille. Efficient multilevel brain tumor segmentation with integrated bayesian model classification. *IEEE Transactions on Medical Imaging*, 27(5):629–640, 2008a.

J. J. Corso, Z. Tu, and A. Yuille. MRF Labeling with a Graph-Shifts Algorithm. In *Proceedings of International Workshop on Combinatorial Image Analysis*, volume LNCS 4958, pages 172–184, 2008b. 3

J. M. Coughlan and A. L. Yuille. Algorithms from Statistical Physics for Generative Models of Images. *Image and Vision Computing, Special Issue on Generative-Model Based Vision*, 21(1):29–36, 2003. 3

H. Derin and H. Elliott. Modeling and segmentation of noisy and texture images using gibbs random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(1):39–55, 1987. 3

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

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

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

T. S. Ferguson. A bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1(2):209–230, 1973. 3

D. J. Field. 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, 1987. 2

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

D. Forsyth, J. Haddon, and S. Ioffe. The joy of sampling. *International Journal of Computer Vision*, 41(1):109–134, 2001. 3

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

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

B. Gidas. A Renormalization Group Approach to Image Processing Problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(2):164–180, 1989. ISSN 0162-8828. doi: 10.1109/34.16712. 3

P. J. Green. Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika*, 82(4):711–732, 1995. 3

T. L. Griffiths and A. Yuille. Technical introduction: A primer on probabilistic inference. Technical report, University of California at Los Angeles, 2006. 2

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

C. E. Guo, S. C. Zhu, and Y. N. Wu. Primal sketch: Integrating texture and structure. *Computer Vision and Image Understanding*, 2006. (to appear). 3

F. Han and S. C. Zhu. Bottom-up/top-down image parsing by attribute graph grammar. In *Proceedings of International Conference on Computer Vision*, volume 2, pages 1778–1785, 2005. 3

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

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

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

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

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

P. Kohli, M. P. Kumar, and P. H. S. Torr. P^3 & beyond: Solving energies with higher order cliques. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2007. 3

P. Kohli, L. Ladicky, and P. H. S. Torr. Robust higher order potentials for enforcing label consistency. *International Journal of Computer Vision*, 82:302–324, 2009. 3

V. Kolmogorov and R. Zabih. What Energy Functions Can Be Minimized via Graph Cuts? In *European Conference on Computer Vision*, volume 3, pages 65–81, 2002a.

V. Kolmogorov and R. Zabih. What Energy Functions Can Be Minimized via Graph Cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, 2004. 3

Vladimir Kolmogorov and Ramin Zabih. Multicamera Scene Reconstruction via Graph-Cuts. In *European Conference on Computer Vision*, pages 82–96, 2002b.

S. Krishnamachari and R. Chellappa. Multiresolution gmrf models for texture segmentation. volume 4, pages 2407–2410 vol.4, 1995. 3

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

L. Ladicky, C. Russell, P. Kohli, and P. H. S. Torr. Associative hierarchical crfs for object class image segmentation. In *Proceedings of International Conference on Computer Vision*, 2009. 3

J. Lafferty, A. McCallum, and F. Pereira. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of International Conference on Machine Learning*, pages 282–289, 2001. 3

C. H. Lee, M. Schmidt, A. Murtha, A. Bistritz, J. Sander, and R. Greiner. 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*, pages 469–478, 2005.

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

S. Z. Li. *Markov Random Field Modeling in Image Analysis*. Springer-Verlag, 2nd edition, 2001. 2, 3

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

J. S. Liu and R. Chen. Sequential monte carlo methods for dynamic systems. *Journal of the American Statistical Society*, 93(443): 1032–1044, 1998. 3

S. N. MacEachern and P. Muller. Estimating mixture of dirichlet process models. *Journal of Computational and Graphical Statistics*, 7 (2):223–238, 1998.

J. Malik, S. Belongie, J. Shi, and T. Leung. Textons, Contours, and Regions: Cue Combination in Image Segmentation. In *International Conference on Computer Vision*, 1999. 3

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

B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23): 3311–3325, 1997. 3

A. Raj and R. Zabih. A graph cut algorithm for generalized image deconvolution. In *Proceedings of International Conference on Computer Vision*, 2005.

A. Ranganathan. The dirichlet process mixture (dpm) model. September 2004. URL `http://www.cs.rochester.edu/~michalak/mlseminar/fall05/dirichlet.pdf`.

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

S. Roth and M. J. Black. Fields of experts. *International Journal of Computer Vision*, 82(2):205–229, 2009. 3

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

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

E. P. Simoncelli and B. A. Olshausen. Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24:1193–1216, 2001. 2

M. Steyvers, P. Smyth, M. Rosen-Zvi, and T. Griffiths. Probabilistic author-topic models for information discovery. In *10th ACM SigKDD Conference on Knowledge Discovery and Data Mining*, 2004. 3

E. B. Sudderth, A. Torralba, W. T. Freeman, and A. S. Willsky. Describing visual scenes using transformed dirichlet processes. In *Proceedings of Neural Information Processing Systems (NIPS)*, 2005.

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

Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Hierarchical dirichlet processes. In *Advances in Neural Information Processing Systems (NIPS) 17*, 2005. 3

L. Tierney. Markov chains for exploring posterior distributions. *The Annals of Statistics*, 22(4):1701–1728, 1994. 3

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

A. Torralba and A. Oliva. Statistics of natural image categories. *Network: Computation in Neural Systems*, 14:391–412, 2003. 2

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

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

Z. Tu, X. R. Chen, A. L. Yuille, and S. C. Zhu. Image Parsing: Unifying Segmentation, Detection and Recognition. *International Journal of Computer Vision*, 63(2):113–140, 2005. 3

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

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

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

J. S. Yedidia, W. T. Freeman, and Y. Weiss. Generalized belief propagation. In *Advances in Neural Information Processing Systems (NIPS)*, volume 13, pages 689–695, 2000. 3

J. S. Yedidia, W. T. Freeman, and Y. Weiss. Bethe free energy, Kikuchi approximations and belief propagation algorithms. Technical Report TR2001-16, Mitsubishi Electronic Research Laboratories, May 2001. 3

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

Ramin Zabih and Vladimir Kolmogorov. Spatially Coherent Clustering Using Graph Cuts. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 437–444, 2004.

Y. Zhang, M. Brady, and S. Smith. 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, January 2001. 3

S. C. Zhu. Stochastic jump-diffusion process for computing medial axes in markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(11):1158–1169, 1999.

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

S. C. Zhu and D. Mumford. Prior learning and gibbs reaction-diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(11):1236–1250, 1997.

S. C. Zhu and D. Mumford. A stochastic grammar of images. *Foundations and Trends in Computer Graphics and Vision*, 2(4):259–362, 2007. 3

S. C. Zhu and A. Yuille. 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, 1996a. 3

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

S. C. Zhu, Y. Wu, and D. Mumford. Minimax entropy principle and its application to texture modeling. *Neural Computation*, 9(8): 1627–1660, 1997. 3

S. C. Zhu, Y. N. Wu, and D. B. Mumford. FRAME: Filters, Random field And Maximum Entropy: — Towards a Unified Theory for Texture Modeling. *International Journal of Computer Vision*, 27(2):1–20, 1998. 3

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

General Notes

If you don't understand something covered in class, ask about it right away. The only silly question is the one which is not asked. If you get a poor mark on an assignment or exam, find out why right away. Don't wait a month before asking. The instructor and teaching assistant are available to answer your questions. Don't be afraid to ask questions, or to approach the instructor or TA in class, during office hours, through the newsgroup or through e-mail. This course is intended to be hard work, but it is also intended to be interesting and fun. We think pattern recognition is interesting and exciting, and we want to convince you of this.

Disabilities

If you have a diagnosed disability (physical, learning, or psychological) that will make it difficult for you to carry out the course work as outlined, or that requires accommodations such as recruiting note-takers, readers, or extended time on exams or assignments, you must consult with the Office of Disability Services (25 Capen Hall, Tel: 645-2608, TTY: 645-2616, Fax: 645-3116, `http://www.student-affairs.buffalo.edu/ods/`). You must advise your instructor during the rst two weeks of the course so that we may review possible arrangements for reasonable accommodations.

Counseling Center

Your attention is called to the Counseling Center (645-2720), 120 Richmond Quad. The Counseling Center staff are trained to help you deal with a wide range of issues, including how to study effectively and how to deal with exam-related stress. Services are free and condential. Their web site is `http://www.student-affairs.buffalo.edu/shs/ccenter/`.

Distractions In The Classroom - Behavioral Expectations

The following is the text of a policy adopted by the Faculty Senate on 5/2/2000. You are expected to know and adhere to this policy.

OBSTRUCTION OR DISRUPTION IN THE CLASSROOM - POLICIES
UNIVERSITY AT BUFFALO

To prevent and respond to distracting behavior faculty should clarify standards for the conduct of class, either in the syllabus, or by referencing the expectations cited in the Student Conduct Regulations. Classroom "etiquette" expectations should include:

- Attending classes and paying attention. Do not ask an instructor in class to go over material you missed by skipping a class or not concentrating.
- Not coming to class late or leaving early. If you must enter a class late, do so quietly and do not disrupt the class by walking between the class and the instructor. Do not leave class unless it is an absolute necessity.
- Not talking with other classmates while the instructor or another student is speaking. If you have a question or a comment, please raise your hand, rather than starting a conversation about it with your neighbor.
- Showing respect and concern for others by not monopolizing class discussion. Allow others time to give their input and ask questions. Do not stray from the topic of class discussion.
- Not eating and drinking during class time.
- Turning off the electronics: cell phones, pagers, and beeper watches.
- Avoiding audible and visible signs of restlessness. These are both rude and disruptive to the rest of the class.
- Focusing on class material during class time. Sleeping, talking to others, doing work for another class, reading the newspaper, checking email, and exploring the internet are unacceptable and can be disruptive.
- Not packing bookbags or backpacks to leave until the instructor has dismissed class.

Academic Integrity

A zero-tolerance policy on cheating will be adopted in this course. The following is the formal statement of academic integrity. Source: `http://www.cse.buffalo.edu/graduate/policies_acad_integrity.php`

The academic degrees and the research findings produced by our Department are worth no more than the integrity of the process by which they are gained. If we do not maintain reliably high standards of ethics and integrity in our work and our relationships, we have nothing of value to offer one another or to offer the larger community outside this Department, whether potential employers or fellow scholars.

For this reason, the principles of Academic Integrity have priority over every other consideration in every aspect of our departmental life, and we will defend these principles vigorously. It is essential that every student be fully aware of these principles, what the procedures are by which possible violations are investigated and adjudicated, and what the punishments for these violations are. Wherever they are suspected, potential violations will be investigated and determinations of fact sought. In short, breaches of Academic Integrity will not be tolerated.

University Statements on Academic Integrity

The University at Buffalo Department of Computer Science and Engineering endorses and adheres to the University policy on Academic Integrity. Students should be familiar with that policy, as expressed in the following documents.:

- UB Office of Judicial Affairs statement on Academic Dishonesty. `http://www.ub-judiciary.buffalo.edu/art3a.shtml#integrity`
- UB Undergraduate Catalog statement on Academic Integrity. `http://undergrad-catalog.buffalo.edu/policies/course/integrity.shtml`

Departmental Statement on Academic Integrity in Coding Assignments and Projects

The following statement further describes the specific application of these general principles to a common context in the CSE Department environment, the production of source code for project and homework assignments. It should be thoroughly understood before undertaking any cooperative activities or using any other sources in such contexts.

All academic work must be your own. Plagiarism, defined as copying or receiving materials from a source or sources and submitting this material as one's own without acknowledging the particular debts to the source (quotations, paraphrases, basic ideas), or otherwise representing the work of another as one's own, is never allowed. Collaboration, usually evidenced by unjustifiable similarity, is never permitted in individual assignments. Any submitted academic work may be subject to screening by software programs designed to detect evidence of plagiarism or collaboration.

It is your responsibility to maintain the security of your computer accounts and your written work. Do not share passwords with anyone, nor write your password down where it may be seen by others. Do not change permissions to allow others to read your course directories and files. Do not walk away from a workstation without logging out. These are your responsibilities. In groups that collaborate inappropriately, it may be impossible to determine who has offered work to others in the group, who has received work, and who may have inadvertently made their work available to the others by failure to maintain adequate personal security. In such cases, all will be held equally liable.

These policies and interpretations may be augmented by individual instructors for their courses. Always check the handouts and web pages of your course and section for additional guidelines.

Departmental Policy on Violations of Academic Integrity

Any student accused of a violation of academic integrity will be so notified by the course director. An informal review will be conducted, including a meeting between these parties. After this review and upon determination that a violation has occurred, the following sanctions will be imposed. **It is the policy of this department that, in general, any violation of academic integrity will result in an F for the course, that all departmental financial support including teaching assistantship, research assistantship or scholarships be terminated, that notification of this action be placed in the student's confidential departmental record, and that the student be permanently ineligible for future departmental financial support.** A second violation of academic integrity will cause the department to seek permanent dismissal from the major and bar from enrollment in any departmental courses. Especially flagrant violations will be considered under formal review proceedings, which may in addition to the above sanctions result in expulsion from the University.