READING LIST
ARTIFICIAL INTELLIGENCE QUALIFYING EXAMINATION

Fall 2001

1. INTRODUCTION

The *AI Depth Exam* is designed to test for both a general "breadth" knowledge of AI, plus a deeper specialized knowledge of one particular sub-area within AI. The current sub-areas tested for, and the corresponding courses, are: *machine learning* (CS 760), and *computer vision* (CS 766). NOTE: *robot motion planning* (CS 780) IS NO LONGER ONE OF THE SUB-AREAS COVERED BY THE AI QUALIFYING EXAMINATION.

Students will be required to specify their sub-area when they sign up for the exam, and then answer corresponding questions in that sub-area, as described in detail below. Passing the exam will then be considered as demonstration of competence in that particular sub-area, so that except in unusual circumstances, the student will be expected to pursue Ph.D. research in the area chosen.

In the reading list, the General Readings cover the basic material in CS 540. For the Specialized Readings, the readings for each course (except CS 731 and 776) are split into Basic Topics and Advanced Topics.

Students will be prepared for the AI Exam if they have mastered (a) the material presented in CS 540 and in at least two of CS 731, 760, 766, and 776; (b) the General AI Readings (CS 540); the Basic and the Advanced Readings associated with one of CS 760 and 766; and the Basic Readings for at least one additional class chosen from CS 731, 760, 766, and 776. Students should note that the Advanced Readings go beyond the material covered in the associated course.

The AI Examination will consist of questions in six separate sections. There will be two "general" AI questions (in section G540), two "basic" questions in each of sections B731, B760, B766, B776, and two "advanced" questions in each of sections A760 and A766. To pass the examination, students will have to answer satisfactorily one of the two questions in the section G540; both of the questions in one of the A sections (A760 or A766); both of the questions in the B section (B760 or B766) for the same sub-area; and two additional questions from any of the other B sections. The two additional questions can be but do not have to be from the same section. The A section chosen must correspond to the sub-area for which the student is signed up.

The advanced questions focus on principles and concepts, and the student is expected to be able to solve new and original problems which require adapting the material covered in the reading lists below. The seven questions answered by the student will not in general be weighted equally. In particular, students will not pass unless they have demonstrated competence in their chosen specialty.

Copies of many of the papers should be available on the web pages of the authors (use a search engine). Others can be obtained from Wendt Library.

2. GENERAL READINGS

2.1. Topics

1. *Search:* State-space search, heuristic methods, problem-reduction methods, constraint propagation, alpha-beta, simulated annealing, genetic algorithms.
2. *Knowledge Representation and Reasoning:* First-order predicate calculus, semantic networks, neural networks, Bayesian networks. The physical symbol hypothesis. Forward and backward chaining, production systems. Resolution theorem proving.
3. *Planning:* State-space, problem reduction, conjunctive, partial-order, reactive, hierarchical. Situation calculus. The frame problem. STRIPS.
4. *Machine Learning:* Supervised learning. Decision tree induction. Backpropagation training of neural networks.
5. *Natural Language Understanding:* Syntax and semantics. Statistical natural language processing.
6. *Perception:* Edge detection (using convolution) and line detection (using the Hough transform) for computer vision. Hidden Markov models for speech recognition.

2.2. References

The following texts are possible sources for the above topics.

[DEAN95] Dean, T., Allen, J., and Aloimonos, Y., *Artificial Intelligence: Theory and Practice*, Benjamin/Cummings, Redwood City, CA, 1995.

[LUG93] Luger, G. F., and Stubblefield, W. A., *Artificial Intelligence*, 3rd ed., Benjamin/Cummings, Redwood City, CA, 1998.

[RICH91] Rich, E., and Knight, K., *Artificial Intelligence*, 2nd ed., McGraw-Hill, New York, 1991.

[RUSS95] Russell, S., and Norvig, P., *Artificial Intelligence: A Modern Approach*, Prentice Hall, Englewood Cliffs, NJ, 1995.

[WINST92] Winston, P., *Artificial Intelligence*, 3rd ed., Addison-Wesley, Reading, MA, 1992.

3. SPECIALIZED READINGS

3.1. CS 731 — ADVANCED ARTIFICIAL INTELLIGENCE

3.1.1. Basic Topics

1. *Uncertainty and Probability:* Independence, conditional independence, Bayes theorem, normalization, chain rule, product rule. [RUSS00, Chapter 15]
2. *Bayesian Networks:* Syntax, semantics, and basic inference. Definition of a Bayes net, meaning of a Bayes Net (as a full joint), inference by enumeration, variable elimination. Junction trees, moralization, triangulation, upward and downward passes, relationship to variable elimination. Learning in Bayesian networks. [DAMB 99; RUSS00, Chapter 16 and 19; HUAN94]
3. *Approximate Inference in Bayesian Networks:* Inference by Markov Chain Monte Carlo, Gibbs sampling, Metropolis-Hastings algorithm. [GILK95, Chapter 1; RUSS00, Chapter 16]

4 *Learning in Bayesian Networks:* Beta and Dirichlet distributions, EM and Gibbs sampling for learning. [HECK96]

5. *Inductive Logic Programming:* Relational learning, FOIL, case studies. [FINN98; QUIN90]
6. *Text Mining:* Information retrieval. Vector space model, TFIDF, stemming. Text classification and information extraction. [CRAV00; MANN99, pp. 529-541; MITC97, Sections 6.1-3, 6.9, and 6.10]
7. *Planning:* Advanced methods for creating plans. GRAPHPLAN and SATPLAN. Partially Observable Markov Decision Processes (POMDPs) and dynamic Bayesian networks. [RUSS00, Chapter 17; WELD99]

3.1.2. References

[CRAV00] Craven, M. and Slattery, S. , Relational learning with statistical predicate invention: Better models for hypertext, *Machine Learning*, **43(1-2),** 2001, 97-119.

[DAMB99] d'Ambrosio, B., Inference in Bayesian networks, *AI Magazine* **20(2),** 1999, 21-36. (http://www.aaai.org/Library/Magazine/Vol20/vol20.html)

[FINN98] Finn, P., Muggleton, S., Page, D. and Srinivasan, A., Pharmacophore discovery using the inductive logic programming system PROGOL, *Machine Learning* **30(2/3),** 1998, 241-270.

[GILK95] Gilks, W., Richardson, S. and Spiegelhalter, D., *Markov Chain Monte Carlo in Practice.* CRC Press, 1995, 517.

[HECK96] Heckerman, D., *A Tutorial on Learning With Bayesian Networks.* Technical Report MS-TR-95-06 (Revised November 1996).

[HUAN94] Huang, C. and Darwiche, A., Inference in Belief Networks: A Procedural Guide. *International Journal of Approximate Reasoning* **11** 1994.

[MANN99] Manning, C. and Schutze, H., Chapter 15, Topics in Information Retrieval, *Foundations of Statistical Natural Language Processing*, 1999, MIT Press.

[MITC97] Mitchell, T., *Machine Learning*, McGraw-Hill, New York, 1997.

[QUIN90] Quinlan, R., Learning logical definitions from relations, *Machine Learning* **5,** 1990, 239-266.

[RUSS00] Russell, S. and Norvig, P., Artificial Intelligence: A Modern Approach. **2nd**Edition, (chapter copies available from D. Page).

[WELD99] Weld, D., Recent advances in AI planning, *AI Magazine* **20(2),** 1999, 93-123. (http://www.aaai.org/Library/Magazine/Vol20/vol20.html)

3.2. CS 760 — MACHINE LEARNING

3.2.1. Basic Topics

1. *Empirical Learning:* Inductive learning, inductive bias, handling noisy and incomplete data. Overfitting. ID3, Version spaces, Naive Bayes. Ensemble methods. [DIET97a, section on ensembles; MITC97, Chapters 2 and 3, Section 6.9]
2. *Connectionism and Genetic Algorithms:* Perceptrons, delta rule, backpropagation, hidden units. Comparisons to symbolic approaches for empirical learning. Genetic algorithms. [MITC97, Chapters 5 and 9]
3. *Learning by Observation:* Conceptual clustering. COBWEB. [FISH87]
4. *Explanation-Based Learning:* Use of domain-specific knowledge. Operationality vs. generality, imperfect domains, the utility problem. [MITC97, Chapter 11]
5. *Case-Based Learning:* Instance-based models. [MITC97, Chapter 8]
6. *Inductive Logic Programming:* Relational learning. FOIL. [MIT97, Chapter 10]
7. *Reinforcement Learning:* Embedded learning systems. Q-learning. [MITC97, Chapter 13]
8. *Experimental Methodology:* Training/tuning/testing sets, *N*-fold cross-validation, ROC curves, *t*-tests. [MITC97, Chapter 5; PROV98]

3.2.2. Advanced Topics

1. *General Issues in Empirical Learning:* Knowledge-level learning, constructive induction, minimal description length (MDL) methods, the bias-variance dilemma, ensemble methods, boosting and bagging, co-training, multi-instance learning, support-vector machines. Experimental methodology. [BLUM98; DIET86; DIET97a; DIET97b; DIET98; DOMI99; FREU96; GEMA92; HEAR98; HECK96; MITC97, Chapters 6 and 10; BREI96; QUIN89; QUIN90; QUIN96]
2. *Neural Networks:* Training neural networks. Weight decay, radial basis functions, recurrent networks. Deciding what expression should be optimized. Rule extraction from trained networks. Connectionist unsupervised learning. [CRAV96; ELMA90; MITC97, Chapter 4 and Section 8.4; RUME85; RUME95]
3. *Reinforcement Learning:* Temporal difference learning, exploration-exploitation tradeoff, Q-learning. SARSA. [BART95; DIET97a; KAEL96; SUTT88]
4. *Inductive Logic Programming:* Relational learning. Top-down and bottom-up induction. FOIL, FOCL, and GOLEM. [MUGG92; QUIN90; PAZZ92]

5. *Using Prior Knowledge:* Using both existing knowledge and training examples to learn. Theory refinement, learning in Bayesian networks. [HECK96; PAZZ92; TOWE94]

3.2.3. References

[BART95] Barto, A., Bradtke, J., and Singh, S., Learning to act using real-time dynamic programming, *Artificial Intelligence* **72**, 1995, 81-138.

[BLUM98] Blum, A. and Mitchell, T., Combining labeled and unlabeled data with co-training, *Proc. 11th Annual Conf. on Computational Learning Theory*, 1998, 92-100.

[BREI96] Breiman, L., Bagging predictors, *Machine Learning* **24**, 1996, 123-140.

[CRAV96] Craven, M. and Shavlik, J., Extracting tree-structured representations of trained networks, *Advances in Neural Information Processing Systems*, Volume 8, MIT Press, 1996 (ftp://ftp.cs.wisc.edu/machine-learning/shavlik-group/craven.nips96.ps)

[DIET86] Dietterich, T., Learning at the knowledge level, *Machine Learning* **1**, 1986, 287-316.

[DIET97a] Dietterich, T., Machine learning research: Four current directions, *AI Magazine* **18**, 1997, 97-136. (http://www.aaai.org/Library/Magazine/Vol18/vol18.html)

[DIET97b] Dietterich, T., Lathrop, R., nd Lozano-Perez, T, Solving the multiple-instance problem with axis-parallel rectangles, *Artificial Intelligence*, **89**, 1997, 31-71.

[DIET98] Dietterich, T., Approximate statistical tests for comparing supervised classification learning algorithms, *Neural Computation*, **10**, 1998, 1895-1924.

[DOMI99] Domingos, P., The role of Occam's Razor in knowledge discovery. *Data Mining and Knowledge Discovery*, **3** (4), 1999.

[ELMA90] Elman, J., Finding structure in time, *Cognitive Science* **14**, 1990, 179-211.

[FISH87] Fisher, D., Knowledge acquisition via incremental conceptual clustering, *Machine Learning* **2**, 1987, 139-172.

[FREU96] Freund, Y. and Schapire, R., Experiments with a new boosting algorithm, in *Proc. 13th Intl. Conf. on Machine Learning*, 1996, 148-156.

[GEMA92] Geman, S., Bienenstock, E., and Doursat, R. Neural networks and the bias/variance dilemma, *Neural Computation* **4**, 1992, 1-58

[HEAR98] Hearst, M. (editor), Support Vector Machines, *IEEE Intelligent Systems* **13**, 1998, 18-28.

[HECK96] Heckerman, D., *A Tutorial on Learning With Bayesian Networks.* Technical Report MS-TR-95-06 (Revised November 1996). (ftp://ftp.research.microsoft.com/pub/tr/tr-95-06.ps)

[KAEL96] Kaelbling, L., Littman, M., and Moore, A., Reinforcement learning: A survey, *J. of Artificial Intelligence Research*, **4**, 1996, 237-285 (visit www.jair.org).

[MITC97] Mitchell, T., *Machine Learning*, McGraw-Hill, New York, 1997.

[MUGG92] Muggleton, S. and Feng, C., Efficient induction of logic programs, in *Inductive Logic Programming* (Chapter 13), S. Muggleton (Ed.), Academic Press, 1992, 281-298.

[PAZZ92] Pazzani, M. and Kibler, D., The utility of knowledge in inductive learning, *Machine Learning* **9**, 1992, 57-94.

[PROV98] Provost, F., Fawcett, To., and Kohavi, R., The case against accuracy estimation for comparing induction algorithms, *Proc. 15th Intl. Conf. on Machine Learning*, 1998, 445-453.

[QUIN89] Quinlan, J. R. and Rivest, R., Inferring decision trees using the minimal description length principle, *Information and Computation* **80**, 1989, 227-248.

[QUIN90] Quinlan, J. R., Learning logical definitions from relations, *Machine Learning* **5**, 1990, 239-266.

[QUIN96] Quinlan, J. R., Bagging, boosting, and C4.5, *Procs. of the 13th National Conference on Artificial Intelligence*, 1996, 725-730.

[RUME85] Rumelhart, D. and Zipser, D., Competitive learning, *Cognitive Science* **9**, 1985, 75-112.

[RUM95] Rumelhart, D., Durbin, R., Golden, R., and Chauvin, Y., Backpropagation: The basic theory, in *Backpropagation: Theory, Architectures, and Applications*, Y. Chauvin and D. E. Rumelhart (eds.), pp. 1-34, Erlbaum, 1995.

[SUTT88] Sutton, R. Learning to predict by the methods of temporal differences, *Machine Learning* **3**, 1988, 9-44.

[TOWE94] Towell, G. and Shavlik, J., Knowledge-based neural networks, *Artificial Intelligence* **70**, 1994, 119-165. (ftp://ftp.cs.wisc.edu/machine-learning/shavlik-group/towell.aij94.ps)

3.3. CS 766 — COMPUTER VISION

3.3.1. Basic Topics

1. *Edge Detection:* First and second differential operators, Laplacian-of-Gaussian operator, Canny operator, Gaussian and Laplacian pyramids. [BURT83; TRUC98 (Chap. 3, 4); FORS01 (Chap. 9)]
2. *Segmentation:* Thresholding, boundary and curve detection, chain code, Hough transform, RANSAC algorithm. [FISC81; TRUC98 (Chap. 5.1-3); FORS01 (Chap. 16.3-4, 17.1, 18.2.3)]
3. *Inferring Three-Dimensional Shape:* 2½D sketch, stereo. [MARR78; TRUC98 (Chap. 7.1, 7.2); FORS01 (Chap. 12.2, 13)]
4. *Two-Dimensional Object Recognition:* 2D object representations, model-based matching. [TRUC98 (Chap. 10 (except 10.3))]

3.3.2. Advanced Topics

1. *Image Formation, Camera Calibration, and Image Synthesis:* Imaging geometry, radiometry, orthographic and perspective projection, digitization, mosaics, view synthesis. [TRUC98 (Chap. 2, 6); FORS01 (Chap. 1, 2, 6); AVID97; SEIT96; KUTU00; SZEL96]
2. *Feature Detection, Segmentation, and Grouping:* Scale-space, relaxation labeling, snakes, region growing, texture, perceptual organization and grouping, RANSAC. [TRUC98 (Chap. 3-5); FORS01 (Chap. 16)]

3. *Inferring Three-Dimensional Shape:* Epipolar geometry, multiple view geometry, fundamental matrix, trilinear tensor, shape from shading, stereo, photometric stereo, shape from contour, shape from texture. [TRUC98 (Chap. 7, 9); FORS01 (Chap. 3, 11.4, 12); BLOS89; LUCA81; OKUT93; XU96]
4. *Motion Analysis:* Motion detection and optical flow, tracking, snakes, layers, structure from motion, factorization method, spatiotemporal processing. [TRUC98 (Chap. 5.4, 8); FORS01 (Chap. 14); BLAC96; DAVI97; LUCS81; TOMA91; WANG94; BEAR96]
5. *Object Recognition:* 2D and 3D object model construction and representation, appearance-based methods, alignment method, geometric hashing, model-based matching, invariants, viewpoint and pose determination, systems. [TRUC98 (Chap. 10, 11.1, 11.2); FORS01 (Chap. 20); HUTT93; MURA95; WOLF90]
6. *Active Vision:* Attention, focus, viewpoint, vergence, tracking. [KROT87; KUTU94]

3.3.3. References

[AVID97] Avidan, S., and Shashua, A., Novel view synthesis in tensor space, *Proc. Computer Vision and Pattern Recognition Conf.*, 1997, 1034-1040.

[BEAR96] Beardsley, P., Torr, P., and Zisserman, A., 3D model acquisition from extended image sequences, *Proc. 4th European Conf. Computer Vision, Vol. II,* 1996, 683-695.

[BLAC96] Black, M., and Jepson, A., EigenTracking: Robust matching and tracking of articulated objects using a view-based representation, *Proc. 4th European Conf. Computer Vision, Vol. I*, 1996, 329-342.

[BLOS89] Blostein, D., and Ahuja, N., Shape from texture: Integrating texture-element extraction and surface estimation, *IEEE Trans. Pattern Analysis and Machine Intelligence* **11**, 1989, 1233-1251.

[BURT83] Burt, P., and Adelson, E., The Laplacian pyramid as a compact image code, *IEEE Trans. Communications* **31**, 1983, 532-540.

[DAVI97] Davis, J., and Bobick, A., The representation and recognition of action using temporal templates, *Proc. Computer Vision and Pattern Recognition Conf.*, 1997, 928-934.

[FISC81] Fischler, M. and Bolles, R., *Random sample consensus: A paradigm for model fitting with application to image analysis and automated cartography, Comm. ACM* **24**, 1981, 381-395.

[FORS01] Forsyth, D. and Ponce, J., *Computer Vision: A Modern Approach*, Prentice Hall, Upper Saddle River, N.J., 2001. http://www.cs.berkeley.edu/~daf/book.html

[HUTT93] Huttenlocher, D., Recognition by alignment, in *Three-Dimensional Object Recognition Systems*, A. Jain and P. Flynn, eds., Elsevier, Amsterdam, 1993, 311-326.

[KROT87] Krotkov, E., Focusing, *Int. J. Computer Vision* **1**, 1987, 223-237.

[KUTU94] Kutulakos, K., and Dyer, C., Occluding contour detection using affine invariants and purposive viewpoint control, *Proc. Computer Vision and Pattern Recognition Conf.*, 1994, 323-330.

[KUTU00] Kutulakos, K., and Seitz, S., A theory of shape by space carving, *Int. J. Computer Vision* **38**, 2000, 199-218. http://www.cs.washington.edu/homes/seitz/papers/kutuijcv00.pdf

[LUCA81] Lucas, B. and Kanade, T., An iterative image registration technique with an application to stereo vision, *Proc. Int. Joint Conf. Artificial Intelligence*, 1981, 674-679.

[MARR78] Marr, D., and Nishihara, H., Visual information processing: Artificial intelligence and the sensorium of sight, *Technology Review* **81**, 1978, 2-23.

[MURA95] Murase, H., and Nayar, S., Visual learning and recognition of 3-D objects from appearance, *Int. J. Computer Vision* **14**, 1995, 5-24.

[OKUT93] Okutomi, M. and Kanade, T., A multiple baseline stereo, *IEEE Trans. Pattern Analysis and Machine Intelligence* **15**, 1993, 353-363.

[SEIT96] Seitz, S., and Dyer, C., View morphing, *Proc. SIGGRAPH 96*, 1996, 21-30.

[SZEL96] Szeliski, R., Video mosaics for virtual environments, *IEEE Computer Graphics and Applications* **16**(2), 1996, 22-30.

[TOMA91] Tomasi, C., and Kanade, T., Factoring image sequences into shape and motion, *Proc. IEEE Workshop on Visual Motion*, 1991, 21-28. (Or see: Tomasi, C., and Kanade, T., Shape and motion from image streams under orthography: A factorization method, *Int. J. Computer Vision* **9**, 1992, 137-154.)

[TRUC98] Trucco, E., and Verri, A., *Introductory Techniques for 3-D Computer Vision*, Prentice Hall, Upper Saddle River, N.J., 1998.

[WANG94] Wang, J., and Adelson, E., Representing moving images with layers, *IEEE Trans. Image Processing* **3**, 1994, 625-638.

[WOLF90] Wolfson, H., Model-based object recognition by geometric hashing, *Proc. 1st European Conf. Computer Vision*, O. Faugeras, ed., Springer-Verlag, New York, 1990, 526-536.

[XU96] G. Xu and Z. Zhang, *Epipolar Geometry in Stereo, Motion and Object Recognition*, Kluwer, Boston, 1996, Chapter 2.

3.4. CS 776 — ADVANCED BIOINFORMATICS

1. *Unsupervised Learning and Clustering:* Expectation maximization (EM) methods, k-means clustering, hierarchical clustering, graph-based clustering. [DUDA00 chapter 10, MANN chapter 14]
2. *Markov Chain Models:* High-order Markov chains, inhomogeneous Markov chains. [DURB98, chapter 3]
3. *Hidden Markov Models:* Vitberbi, Forward and Backward algorithms, Baum-Welch parameter estimation. [DURB98, chapters 3 and 5]
4. *Stochastic Context-Free Grammars:* Chomsky hierarchy of transformational grammars, stochastic grammars, the Inside, Outside and CYK algorithms, Inside-Outside parameter estimation [DURB98, chapter 9]

3.4.1. References

[DUDA00] Duda, R., Hart, P. and Stork, D., *Pattern Classification*, 2nd edition, John Wiley and Sons, 2000.

[DURB98] Durbin, R., Eddy, S., Krogh, A. and Mitchison, G., *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*, Cambridge University Press, 1998.

[MANN99] Manning, C. and Schutze, H., *Foundations of Statistical Natural Language Processing*, MIT Press, 1999.