

Introductor Applied Machine Learning

School of Informatics, University of Edinburgh

Instructors: Dr. Victor Lavrenko, Dr. Nigel Goddard

September, 2012

Assessment is by formal examination (75%) and four equally-weighted assignments together worth 25%. The assignments will be released in weeks 3, 5, 7 and 9, and will be due at the end of weeks 4, 6, 8 and 10 respectively.

Online material is available at <http://www.inf.ed.ac.uk/teaching/courses/iaml>.

Resources: The course textbook is “Data Mining: Practical Machine Learning Tools and Techniques” (Second Edition) I.H. Witten and E. Frank, Morgan Kaufmann, 2005. ISBN 0-12-088407-0.

Tutorials and Labs: There will be 4 tutorials and 4 labs classes for the course, in weeks 3,5,7 and 9. There will be no labs or tutorials in weeks 1 and 2.

Syllabus Note: this list is provisional

- Introduction (NG)
- Thinking about data (VL)
- Basic probability (NG)
- Naive Bayes classification (VL)
- Nearest-neighbor methods (VL)
- Decision trees (VL)
- K Nearest Neighbours and KD-Trees (VL)
- Linear regression (NG)
- Generalization (NG)
- Logistic regression, perceptrons (NG)
- Support vector machines (NG)
- Dimensionality reduction (VL)
- Clustering (k-means, hierarchical) (VL)
- Mixture models and the EM algorithm (VL)
- Performance evaluation (VL)
- Further topics as time permits, e.g. Mixtures of Experts, Boosting and Bagging, Hidden Markov Models, dealing with Big Data, Collaborative Filtering