Detailed syllabus:

IMPORTANT: this is subject to change depending on how the lectures actually proceed

Lecture 1: fundamentals of probability and Bayes theorem.

Lecture 2: Probability distributions and the Gaussian distribution, maximum likelihood principle.

Lecture 3: Multivariate Gaussian distribution: conditionals and marginals.

Lecture 4: Bayesian computations with Gaussians.

Lecture 5: Linear dimensionality reduction methods: probabilistic PCA.

Lecture 6: Linear regression: maximum likelihood and Bayesian inference.

Lecture 7: Basis functions regression.

Lecture 8: Gaussian processes for regression.

Lecture 9: Generalised linear models (if time permits)

References (all freely available online):

- Information theory, Inference and Learning Algorithms by David Mackay (particularly on probability distributions)

- Bayesian Reasoning and Machine Learning by David Barber

- For lectures 6-8, Gaussian Processes for Machine Learning Ch 2, by Carl Rasmussen and Christopher Williams

- For lecture 5, Probabilistic Principal Component Analysis, by M. Tipping and C. Bishop, J.Roy. Soc. Ser. B, 1998