

Introduction to Probabilistic Graphical Models

The actual science of logic is conversant at present only with things either certain, impossible, or entirely doubtful, none of which (fortunately) we have to reason on. Therefore the *true logic* for this world is the calculus of *Probabilities*, which takes account of the magnitude of the probability which is, or ought to be, in a reasonable man's mind. – James Clerk Maxwell

Study group web site: <http://research.ee.sun.ac.za/pgms/>

Note: Since I'll be updating the instructions as the course proceeds, the updated instructions might only become available directly before/at the time of the meeting.

Organised by: JA du Preez, Phone: x4342, Email: dupreez at sun

Language: Afrikaans or English, depending on participants.

Format: This is intended as a guided self-study, with group support, to get you started in this field. It is NOT a formal offering by US, there is NO formal evaluation and you will NOT be able to get formal credits for this. There will be weekly reading material and exercises (that will involve coding) to do, followed by a discussion at the next session. The work load should be around 5-10 hours per week.

Prerequisites: You will need to be familiar with basic probability theory and should be able to write and use software – the PGM toolbox we use is written in c++.

Content: We will cover the following probabilistic graphical model (PGM) concepts at an introductory level:

Representation: reasoning patterns, Bayes networks (BNs), Markov random fields (MRFs), independence and graphical structure, factor graphs, cluster graphs, junction trees

Inference: variable elimination (VE), belief propagation (BP), belief update (BU), the sum-product algorithm vs. the max-product algorithm

Learning: maximum-likelihood estimation (MLE), maximum a posteriori (MAP) estimation, latent variables and the expectation maximisation (EM) algorithm, introduction to Bayesian learning

Primary resources:

Book: Bayesian Reasoning and Machine Learning, David Barber, Cambridge University Press, 2012. Free pdf download available [here](#).

Video lectures: We will use Daphne Koller's PGM lecture series. The videos are part of the Probabilistic Graphical Models Specialization on Coursera. Within this specialization, there are the following 3 courses:

- Probabilistic Graphical Models 1: Representation
- Probabilistic Graphical Models 2: Inference
- Probabilistic Graphical Models 3: Learning

To access the videos, follow this procedure:

1. Register on Coursera.
2. Go to the website of one of the 3 PGM courses.
3. Click on "Enroll for free".
4. At the bottom of the pop-up window, click on "Audit the course". You should now have access to all the course material, including the videos.

Code: For the coding parts of the assignments, we will use Stellenbosch University's `emdw` toolbox. For this you will need Linux and a fairly new g++ compiler. Setup instructions will be provided separately.

Other resources:

Online course: Duration 10 weeks, you will need to invest about 20 hours per week. See <https://www.coursera.org/> for the next round of it.

Supplementary book 1: Probabilistic Graphical Models: Principles and Techniques, Daphne Koller and Nir Friedman, MIT Press, 2009.

Supplementary book 2: Information Theory, Inference and Learning Algorithms, David JC MacKay, Cambridge University Press, 2003. Free pdf download at <http://www.inference.phy.cam.ac.uk/mackay/itila/book.html>