Types and states: Mixture and hidden Markov models for cognitive science

Ingmar Visser (i.visser@uva.nl)

Department of Developmental Psychology, University of Amsterdam
Weesperplein 4, 1018 XA Amsterdam, The Netherlands

Maarten Speekenbrink (m.speekenbrink@ucl.ac.uk)

Cognitive, Perceptual, and Brain Sciences, University College London
Gower Street, London WC1E 6BT, United Kingdom

Objectives and scope

There are many situations in which one may encounter distinct *types* of entities, such as different animal species, and different *states* in which these entities may exist, for example motivational states like hunger. Cognition is sometimes also best understood in terms of discrete types and states. For example, forms of cognitive development can be characterised as the acquisition of increasingly complex rules which constitute different types of reasoning and associated response patterns in reasoning tasks (Jansen, Raij-makers, & Visser, 2007). And rather than a gradually shifting trade-off, people may switch rapidly between distinct decision-making modes favouring either speed or accuracy (Dutilh, Wagenmakers, Visser, & Maas, 2011). The idea that cognitive processes are guided by qualitatively different strategies underlies a wide range of theories concerned with topics such as word recognition, cognitive development, categorization, and decision making, to name but a few (for an overview, see e.g. Scheibehenne, Rieskamp, & Wagenmakers, 2013).

As the identity of types and states is generally not directly observable, appropriate statistical techniques are required to identify them. This tutorial will focus on mixture models (MMs) and hidden Markov models (HMMs), which are the basis of such techniques. In the context of MMs, a type or state (e.g., a cognitive strategy) is formalized as a probability distribution over observables. Because a dataset may contain different types, the overall distribution is a mixture of such individual component distributions. As the component distributions need not be of the same parametric family (e.g., Gaussian distributions can be mixed with other distributions), MMs allow for considerable flexibility in the definition of types and states. HMMs are a natural extension of MMs, allowing switches between states over time. For example, these models are useful when people can switch between cognitive strategies during a task. In addition to identifying the different states, HMMs allow one to also focus on the process underlying state transitions.

While mixture models (MMs) and hidden Markov models (HMMs) are widely used in fields such as computational biology (e.g., for DNA sequence analysis) and machine learning (e.g., for speech recognition and estimation of topic models), their use in the analysis of cognition and behaviour is rela-tively rare. This is unfortunate, as MMs and HMMs are ideally suited to test and explore important theoretical ideas in cognitive science. The objective of this tutorial is to provide researchers in cognitive science with an accessible introduction to MMs and HMMs and provide them with the necessary skills to apply them in their own research.

Outline of the tutorial

The tutorial is divided into two parts. The first part introduces the theory behind MMs and HMMs. The second part will be more practical, using a number of examples to show (a) how to apply MMs and HMMs with user-friendly and freely available software, (b) how to interpret these models, and (c) how the models can reveal aspects in the data which remain hidden with more traditional analyses. The first part of the tutorial will be delivered as a classroom style lecture. The second part will use a more hands-on approach with practical computer-based examples and exercises. The audience is encouraged to bring a laptop. All necessary software and material will be made available in advance.

Part I: Mixture models

Introduction to mixture models. This part will introduce the basic structure of mixture models and the use of graphical and other techniques to determine whether MMs might be applicable.

Estimation and inference This part will provide an intuitive treatment of maximum likelihood estimation and introduce numerical optimization and Expectation-Maximization (EM), the two main methods for this type of estimation of MMs and HMMs. Practical issues such as starting values and local maxima will also be discussed. We then discuss methods for model selection and how to determine the number of components (i.e., types, or latent classes). We will also discuss methods to test parameters for significance and the use of posterior probabilities to determine the component to which a data point belongs.

Introduction to `depmixS4` This part will introduce `depmixS4` (Visser & Speekenbrink, 2010), a flexible package to estimate MMs and HMMs in the R environment for statistical computing (R Development Core Team, 2010). The examples in the remainder of the tutorial will use this package.

Examples of mixture models Examples will include the use of MMs to detect developmental stages in the liquid conservation task and the use of MMs to detect multiple learning strategies in a category learning task.

Part II: Hidden Markov models

Hidden Markov models. This part will introduce hidden Markov models as a direct extension of mixture models. We will then discuss how to generalize the previously discussed methods of estimation and inference to these models.

Examples of hidden Markov models Examples will include the use of HHMs to analyse speed-accuracy trade-offs and the use of HMMs to model response strategies in multiple cue learning.

Extensions This part will briefly discuss some extensions to basic MMs and HMMs, including the use of covariates to predict the identity of mixture components and states.

About the organizers

The organizers are the developers of `depmixS4` (Visser, Jansen, & Speekenbrink, 2010), a popular `R` package to estimate mixture and hidden Markov models. They are also the authors of an upcoming book on this topic (commissioned by Springer for their "UseR" series) and a recent tutorial on hidden Markov models (Visser, 2011). The organizers have extensive experience in the application of MMs and HMMs to research in developmental and cognitive science (e.g. Speekenbrink, Lagnado, Wilkinson, Jahanshahi, & Shanks, 2010; Visser et al., 2010). They can draw upon this experience to provide the audience with real examples and practical advice relevant to a cognitive science audience.

Justification

Given the success of this tutorial at the previous CogSci conference in Berlin (2013), we offer it again this year.

Theories which propose the existence of distinct types and states are widespread in the cognitive sciences. Traditional statistical analysis, such as t-tests and ANOVAs, or not applicable to test such ideas. MMs have been successfully used to test "toolbox models" of cognition (e.g. Scheibehenne et al., 2013) and HMMs to test discrete strategy switches (e.g. Speekenbrink et al., 2010; Jansen et al., 2007). This tutorial will provide cognitive scientists with the intuitive understanding and practical knowledge of these models necessary to help them apply them to their own research.

Intended audience

This tutorial will be mainly introductory and no specific prior background knowledge is required. While basic knowledge of probability and statistics will be helpful, treatment of the theoretical concepts will largely be conceptual. While familiarity with the `R` environment will be helpful in general, by making the commands available, previous experience is no requirement to follow and replicate the results of the practical examples.

Requirements

Participants would ideally bring a laptop to the tutorial. The required software (`R` and `depmixS4`) is open source and freely and easily obtainable. `R` is available for all major platforms and can be downloaded from `http://www.r-project.org`. The `depmixS4` package can be downloaded from `http://cran.r-project.org/web/packages/depmixS4/` or installed from within `R` through the command `install.packages('depmixS4')`.

Contact details

Ingmar Visser, Department of Developmental Psychology. Weesperplein 4, 1018 XA Amsterdam, The Netherlands. Tel: +31 20 5256723, Email: I.Visser@uva.nl.

Maarten Speekenbrink, Cognitive, Perceptual and Brain Sciences, University College London, Gower Street, London WC1E 6BT, England, UK. Tel: +44 20 7679 8548. Email: m.speekenbrink@ucl.ac.uk

References

Dutilh, G., Wagenmakers, E.-J., Visser, I., & Maas, H. L. J. van der. (2011). A phase transition model for the speed–accuracy trade–off in response time experiments. *Cognitive Science*, *35*, 211-250.

Jansen, B. R. J., Raij-makers, M. E. J., & Visser, I. (2007). Rule transition on the balance scale task: A case study in belief change. *Synthese*, *155*(2), 211-236.

R Development Core Team. (2010). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Available from `http://www.R-project.org` (ISBN 3-900051-07-0)

Scheibehenne, B., Rieskamp, J., & Wagenmakers, E.-J. (2013). Testing adaptive toolbox models: A Bayesian hierarchical approach. *Psychological Review*, *120*, 39–64.

Speekenbrink, M., Lagnado, D. A., Wilkinson, L., Jahanshahi, M., & Shanks, D. R. (2010). Models of probabilistic category learning in Parkinson's disease: Strategy use and the effects of L-dopa. *Journal of Mathematical Psychology*, *54*, 123–136.

Visser, I. (2011). Seven things to remember about hidden Markov models: A tutorial on Markovian models for time series. *Journal of Mathematical Psychology*, *55*, 403—415.

Visser, I., Jansen, B. R. J., & Speekenbrink, M. (2010). A framework for discrete change. In P. C. M. Molenaar & K. M. Newell (Eds.), *Individual pathways of change: Statistical models for analyzing learning and development* (pp. 109–123). Washington: APA.

Visser, I., & Speekenbrink, M. (2010). depmixS4: An R-package for hidden Markov models. *Journal of Statistical Software*, *36*(7), 1–21. Available from `http://www.jstatsoft.org/v36/i07/` (R package, current version available from CRAN)