Inference of dynamic probabilistic internal representations from reaction time data

Balázs Török[1,2], **David G. Nagy**[1,3]
{torok.balazs, nagy.g.david}@wigner.mta.hu
Janacsek Karolina[4,5], **Dezső Németh** [4,5]
{janacsek.karolina, nemethd}@gmail.com
Gergő Orbán[1]
orban.gergo@wigner.mta.hu

Abstract

Sequential predictions are ubiquitous in a learning agent's existence. In order to devise efficient responses in a dynamic environment, one needs to build an internal representation of the latent dynamics of the environment. Humans have been shown to create dynamical models such as intuitive physics that approximate the laws of Newtonian physics and are able to reason about their model in terms of formulating new predictions or imagining hypothetical situations. However, subject-by-subject differences in temporal predictions resulting from variations in subjective internal models and individual learning paths have remained unexplored due to the immense difficulty related to inferring dynamical subjective representations. Cognitive Tomography has been proposed to discover static internal representations from discrete choices. We extend this method in two critical ways: 1, We aim to infer internal representations from a richer set of behavioral measures, specifically we use reaction times; 2, Our goal is to infer a dynamical representation. We demonstrate its utility by predicting reaction times and choices of human participants in a probabilistic learning task on a trial-by-trial basis. Inferred behaviour-based trial-specific subjective predictions can be directly used to test theories of neural underpinnings of computations in physiological and imaging data.

Keywords: reaction time, learning, Bayesian inference

Introduction

Perception has been well characterised by Bayesian ideal observer models where prior beliefs are combined with incoming sensory data according to the rules of Bayesian inference to discover the latent causes underlying observations. In most of these cases, group-level behaviour is well accounted for by models that have generic prior beliefs. However, in complex novel tasks, the momentary representation is expected to vary greatly among individuals and the subject-averaged

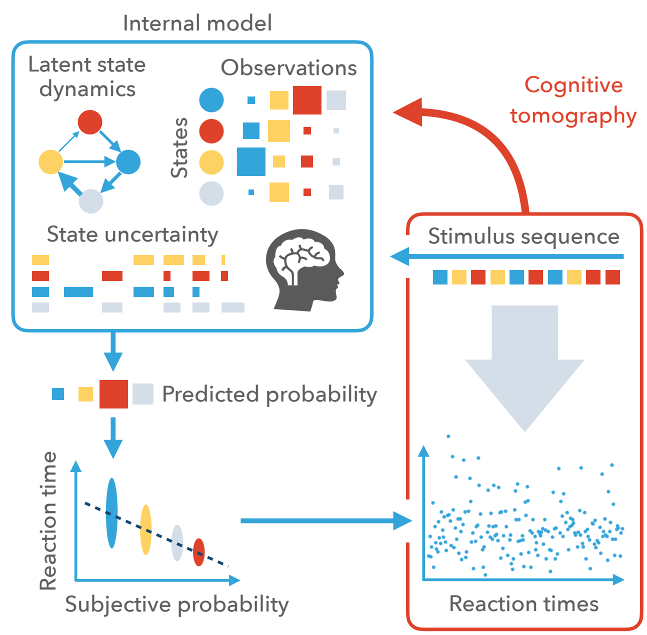

Figure 1: Ideal observer model and Cognitive Tomography. The participant assumes a model for the latent dynamics and a model relating their observations to the latent states. They use these components to update their beliefs over the current state of the observed system and then generate predictions for the upcoming stimulus. The median and standard deviation of reaction times decreases linearly with log subjective probability (R. Carpenter & Williams, 1995). Cognitive Tomography is the method of inverting this generative model. We inferred the internal representations of individuals from the stimulus sequences and the reaction times.

model may be of different form than any of those of the individuals. Therefore, in order to achieve a detailed understanding of learning in a complex environment, we need to extract individual representations with high fidelity.

Cognitive tomography (Houlsby et al., 2013) was recently proposed for inferring multidimensional, detailed individual representations from behavioural data. In its original form, it utilises binary/ternary response data and infers static individual representations for a well trained class of stimuli, human faces. Here we extend cognitive tomography with two major goals in mind that broadens the scope of the original model: First, reaction times can be directly used, enabling the inclu-

[1]Computational Systems Neuroscience Lab, MTA Wigner Research Centre for Physics, Budapest, Hungary
[2]Department of Cognitive Science, Budapest University of Technology and Economics, Budapest, Hungary
[3]Institute of Physics, Eötvös Loránd University, Budapest, Hungary
[4]Institute of Psychology, Eötvös Loránd University, Budapest, Hungary
[5]Institute of Cognitive Neuroscience and Psychology, Hungarian Academy of Sciences, Budapest, Hungary

sion of substantially richer behavioral data; Second, a method for inferring dynamical and potentially evolving internal models is developed. The developed method retains the capacity to infer individualised internal representations and to infer an internal model that is not directly linked to the task the model was trained on. We validate our approach by predicting reaction times on a trial-by-trial level. Critically, while our model is fitted solely on trials with reaction times in correctly executed trials, the inferred internal models also predict the mistakes of participants above chance. Our method provides a new way to assess the acquisition of internal models in a dynamical setting on an individual level.

Ideal observer model

Learning a complex dynamical model is a challenging task. In principle, if one knows the latent dynamics of the system and how their observations relate to the latent dynamics, they can use a filtering technique to track the changes in the latent state of the system and formulate predictions (e.g. Kalman-filters). However, in general one needs to infer the latent dynamical model as well. In a discrete-state world, an exact solution can be given by a Hidden Markov Model (HMM) with a given (possibly infinite) number of states. Recently, Glaze, S Filipowicz, Kable, Balasubramanian, and Gold (2018) demonstrated that a two-state hidden Markov model is an adequate model for internal representation in a simple task. In a general setting, however, the number of states that govern the dynamics of the observations needs to be discovered by the observer and therefore a flexible probabilistic model is required, in which the latent states are learned from data. We use a non-parametric version of HMM, the infinite Hidden Markov Model (iHMM) (Gael, Saatci, Teh, & Ghahramani, 2008), which has been demonstrated to be effective under these circumstances.

An internal model entertained by an ideal observer can be characterized by three factors: the assumed latent dynamics, the generative model of their observations and the momentary beliefs about the current state of the system (Fig 1. Internal model). To link subjective beliefs with behavioral responses a further component is required. An important work by R. Carpenter and Williams (1995) established that saccadic reaction times for correct trials in a probabilistic viewing task were reciprocal normally distributed with fixed variance and median linearly related to the negative log subjective probability (Fig. 1 bottom left). Harris, Waddington, Biscione, and Manzi (2014) argue that manual choice reaction times may also be modelled using reciprocal normal distribution. Our cognitive tomography model uses this generative model to infer an internal representation (and its dimension) from the presented stimulus sequence and the measured reaction times (Fig. 1. Cognitive Tomography).

Method

Experiment

We use the Alternating Serial Reaction Time Task (ASRT) to test if a dynamical internal model can be inferred from reaction time data. In this paradigm, subjects have to manually respond to a stimulus (dog) appearing in one of four locations. The sequence of locations have a structure: in even trials, the stimulus follows a predetermined sequence, while in odd trials, the stimulus appears in any of the four locations with equal probability and independently of all other trials. The statistical structure is unknown to participants and they are instructed to respond to the stimulus as fast and as accurately as possible. After the experiment they reported no explicit knowledge about the statistical structure. The trials are organised into blocks of 80 and 25 blocks are administered with 21 individual participants. The task is well described by an 8-state Hidden Markov Model where there is a non-stochastic latent dynamics repeating the sequence: Pattern1, Random1, Pattern2, Random2, Pattern3, Random3, Pattern4, Random4.

Inference of Internal Model

There are two major challenges we aim to address. First, the model describing the observer's model of the sequence has to be a flexible model able to solve general sequential prediction problems. In theory, any discrete-time sequence can be modelled using Hidden Markov Models. The major challenge is that the number of latent states in such a model needs to be inferred as well. Recent advances in machine learning allow for efficient solution of this problem (Gael et al., 2008), being able to simultaneously infer the number of latent states, the latent dynamics as well as the generative model for observations. A crucial feature of their solution is that they assume an infinite number of latent states, finitely many of which appear in a finite data set. This non-parametric approach allows learning arbitrarily complex models with growing data size. The key idea in their solution is to use slicing, that is, in each step, they carve out of the infinite space of possible models a tractable part and do inference on that part. Consequently, once participants formulate an estimate of the model, they can filter their beliefs over the current state of the world using the latent dynamics and the generative model of the observations. As a point of reference, the algorithm described by Gael et al. (2008) can learn the true generative model of the experiment we conducted from around 200 observations.

The second part of the challenge is to infer the complex, potentially high-dimensional internal model of participants from a sequence of stimuli and behavioral responses (noisy reaction times in our specific case). We can formalise their filtered beliefs, models and the way these generate the observed reaction times as a probabilistic program which repeats the following steps: according to their current beliefs, they calculate expected probabilities for the location of the upcoming stimulus, then depending on their parameters, they produce a reaction time to the upcoming stimulus (R. Carpenter & Williams, 1995). Then, they update their beliefs over the system's state by taking into account the currently observed stimulus and projecting the evolution of the system into the next time-point using their latent dynamics model. Hence, we obtain a full generative model of their reaction time sequence. We, researchers, infer their internal models by inverting this generative model. Since the researcher's posterior (over the participant's model) is high-dimensional and highly structured, we used Hamiltonian Monte Carlo implemented in STAN (B. Carpenter et al., 2017). The Hamiltonian Monte Carlo steps are wrapped into a Gibbs-sampler that is sampling the slicing parameter and the rest of the parameters in an alternating fashion.

We tested our algorithm on synthetically generated reaction time data with reaction time parameters based on the results of R. Carpenter and Williams (1995). The algorithm could recover the ground truth subjective probabilities for up to 0.97 correlation.

Results

First, we contrasted the performance of the dynamical model inferred by cognitive tomography with a classical measure developed to assess learning in the ASRT task. Specifically, we compared the amount of variance explained by cognitive tomography and that explained by the so called triplet model, essentially a trigram model of the input sequence (Howard & Howard, 2001; Janacsek & Nemeth, 2012). We split the 25 blocks of trials into five epochs of equal length and inferred the participants' internal models on each of the five epochs.

Reaction time data was fit using data from one epoch and tested on the following epoch (Fig. 2, left). Exact correspondence between subjective probabilities and reaction time yields a correlation of -1. Deviation from this value is not necessarily the result of discrepancy in the fit of the generative model and the actual internal model of the participant, but the fact that reaction times for individual trials are shown instead of trial averages, which also introduces variance. For all but one participant (Fig. 2, right), our model had better correlation coefficient on a test set than that of the triplet (or trigram) model which states that participants make predictions based on the previous two stimuli, by identifying the most frequent stimulus location following the previous two observed locations.

Next, we aimed at testing the inferred dynamical internal model in a substantially different way: our goal was to demonstrate that the internal model is indeed capable of predicting behaviours that are novel to the model. Critically, the model was trained on reaction times of only those trials to which the participant responded correctly. Subsequently, we tested whether our model forms meaningful predictions for participants' choices in trials where they committed mistakes. In those trials, where the location of the stimulus is independently and randomly generated (with equal probability appearing in any of the positions), hence not predictable from previous stimuli, we expect both reaction times and accuracy to be modulated solely by the participant's internal model. Our model predicts lower subjective probabilities for those random trials that the participant eventually missed (Fig. 3, left). Moreover, we can predict the mistakenly pressed button above chance (for all mistakes on all types of trials, Fig. 3. right; binomial test on first column: $p < 0.01$), based on the internal

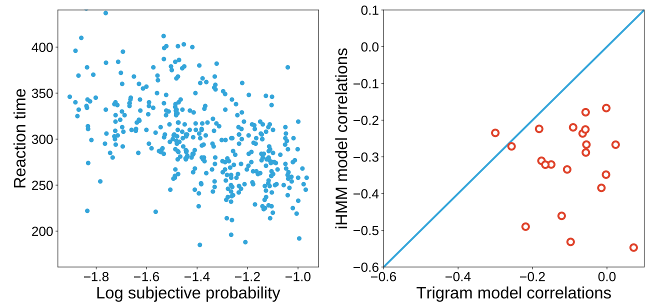

Figure 2: Left: Example reaction times of one participant's last epoch (400 trials) against our model's log predicted probabilities trained on the previous epoch. RTs shown mean ± 3 sd for clarity. Corr: -0.49. Right: Correlations for all individuals' models fitted on the penultimate epoch and correlated with the final epoch. On the x axis, the performance of the model based on trigram predictions. For all but one participant our model has better trial by trial predictions on the test set.

model's predictions. Crucially, we achieve this using no extra parameters, simply by taking the location of the stimulus with largest subjective probability given by the model.

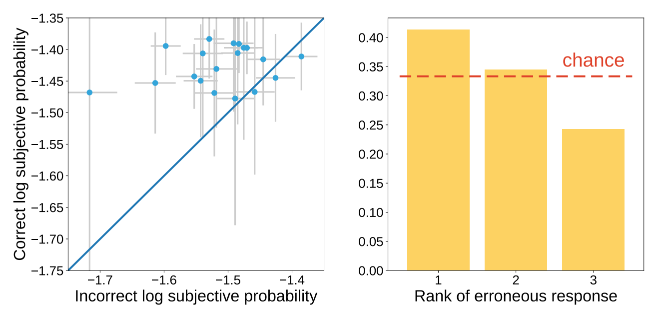

Figure 3: Left: Mean of log subjective probabilities given by the inferred internal model on random trials on a test set. (Model trained on correct trials of the last epoch, predictions are made on missed and correct trials). Dots represent participants with 2 s.e.m. For most participants, the model predicts lower subjective probability to those random (unpredictable) trials which they eventually miss. Right: Ranks of participants' choices on missed trials according to the predictive probabilities of the potential incorrect options. The model predicts participants' actual responses (first column shows the fraction of incorrect choices that had the highest subjective probability according to the internal model) significantly above chance.

Finally, we tested our model at assessing the learning curve of the internal model acquisition. We evaluated how much the internal model changes over time by measuring its performance on different epochs (Fig. 4), which provides an insight into how much the model inferred in one epoch deviates from the models maintained in other epochs. Comparison of the late model (trained on epoch four) with the early

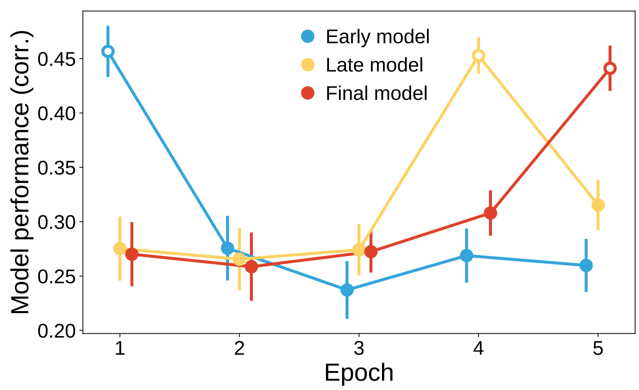

Figure 4: Model performance (negative correlation) of internal models fitted on three different epochs of the experiment (1, 4, 5, for the early, late, final models, respectively). Performance of models decrease with increasing temporal distance from the trained epoch. The model inferred on the first five blocks is inferior to the models inferred at late parts on the late parts of the experiment. Open circles show training performance, filled circles show test performance. Part of the difference between open and filled circles is due to train-test error difference therefore relevant comparisons concern those between filled circles. Error bars show 1 s.e.m. over participants

model (trained on epoch one) reveals that the late model surpasses the early model on the final part of the experiment ($t(20) = -3.98$ $p < 0.01$). Continuity of prediction reveals a gradual change in the internal model. Note, however, that subject-averaged learning curves can hinder the evolution of learning in individuals and tends to provide overly smooth tuning curves. While this analysis promises to reveal learning curves on a subject-by-subject basis, it is the subject of future analysis to reveal individual learning curves.

Conclusions

In this study we extended cognitive tomography to infer a dynamical internal representation and to the use of reaction times. We demonstrated that it can predict reaction times in a probabilistic sequence learning task and surpassed the previously used model by a substantial margin. We also demonstrated that the internal model inferred for individuals can be also used to reason about behaviour not directly trained on, similar to the concept of designing across-task predictions. The model provides subjective probabilities in individual trials and we propose that these explicit trial-by-trial predictions can be used as correlates to physiological or imaging data. The model described here can be easily extended to other behavioural measures, for example eye-movements, which can help investigate acquired internal models in a wide spectrum of tasks.

Acknowledgements

This work has been supported by the National Brain Research Program (project 2017-1.2.1-NKP-2017-00002, PI: D. N.); Hungarian Scientific Research Fund (OTKA PD 124148, PI: K. J.); Janos Bolyai Research Fellowship of the Hungarian Academy of Sciences (K. J.). National Research, Development and Innovation Fund of Hungary (Grant No. K125343, B. T., D. G. N., G. O.) and an MTA Lendület Fellowship (G. O.). Image of head used in Fig. 1. was created by Svelte UX, downloaded from the Noun Project.

References

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software, Articles*, *76*(1), 1–32. doi: 10.18637/jss.v076.i01

Carpenter, R., & Williams, M. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature*, *377*, 59–62.

Gael, J. V., Saatci, Y., Teh, Y. W., & Ghahramani, Z. (2008). Beam Sampling for the Infinite Hidden Markov Model. *Proceedings of the 25th international conference on Machine learning*, 1088–1095.

Glaze, C. M., S Filipowicz, A. L., Kable, J. W., Balasubramanian, V., & Gold, J. I. (2018). A biasâvariance tradeoff governs individual differences in on-line learning in an unpredictable environment. *Nature Human Behaviour*, *2*, 213–224. doi: 10.1038/s41562-018-0297-4

Harris, C. M., Waddington, J., Biscione, V., & Manzi, S. (2014). Manual choice reaction times in the rate-domain. *Frontiers in Human Neuroscience*, *8*. doi: 10.3389/fnhum.2014.00418

Houlsby, N. M. T., Huszar, F., Ghassemi, M. M., Orbán, G., Wolpert, D. M., & Lengyel, M. (2013). Cognitive Tomography Reveals Complex , Task- Independent Mental Representations. *Current Biology*, 2169–2175. doi: 10.1016/j.cub.2013.09.012

Howard, D. V., & Howard, J. H. (2001). When it does hurt to try: Adult age differences in the effects of instructions on implicit pattern learning. *Psychonomic Bulletin & Review*, *8*(4), 798–805.

Janacsek, K., & Nemeth, D. (2012). Predicting the future: From implicit learning to consolidation. *International Journal of Psychophysiology*, *83*(2), 213–221. doi: 10.1016/j.ijpsycho.2011.11.012