Predictive Perception via Simultaneous Learning and Inference

Mahdi Enan[1], **Mario Senden**[1], **Ryszard Auksztulewicz**[1,2], **Federico De Martino**[1]
[1] Maastricht University, Maastricht, The Netherlands, [2] Free University Berlin, Berlin, Germany

Abstract

The brain is continuously faced with noise in the soundscape and uncertainty in the input to the ears. This has led to the proposal that perception is not a direct consequence of sensory input, but rather an inferential process to determine the most probable state of the world. However, the algorithm through which the brain could realize this inferential process as well as its implementation are not yet well understood. In this work, we developed a novel framework for simultaneous learning and inference from first principles using exact inference and local Hebbian-like learning rules. We show that this framework allows unbiased inference and flexible model updating under noise and changing dynamics. Further, we show how this approach can reproduce local and global effects of prediction and prediction error in noisy environments. This model can be used to test (in silico) hypotheses related to predictive processing in noisy and non deterministic environments.

Keywords: perception, inference, Bayes

Introduction

We live in complex noisy environments which cause the input to our brain to be uncertain, an effect that is further amplified by biological limitations in our senses. To interact meaningfully in such uncertain environments, our brains must infer the most likely state of the world from our sensory input. Furthermore, since the dynamics of our environment are not stationary, our brains also have to flexibly adapt to changes in environmental dynamics in order to generate meaningful predictions of upcoming sensory information. Previous work suggested that the brain may employ variational inference motivated by the intractability of exact inference in large state spaces (Friston, 2010). These approaches assume parametrized posterior distributions which, in the discrete case, are approximated via algorithms such as variational message passing. In this work, we argue that exact inference is feasible for small state spaces which is common in experimental settings. Therefore, we present a formal predictive model that captures canonical computational principles underlying our brains' capability to adaptively infer environmental states from noisy sensations. By combining Bayesian filtering with Hebbian-like update rules we provide a framework for exact inference using unsupervised local learning mechanisms. This unified framework allows simultaneous learning and inference by optimally combining sensory information and predictions.

Methods

Our model is composed of a sensor compartment and an internal model of environmental dynamics. The sensor compartment transforms states into a probability distribution of measurements. The sensor can be a simple matrix or a feedforward neural network that maps stimuli (e.g. auditory waveforms) to observations in a measurement space. The internal model keeps track of environmental dynamics in a Markov transition matrix, which is a row-stochastic matrix that upholds the Markov property. At each time-point, the internal model is updated with previous state beliefs using local learning rules that ensure that the transition matrix retains well-formed probability distributions. This state transition matrix is capable of generating predictions of upcoming states by combining this transition matrix with the previous belief(s). A new state belief is formed using the recursive Bayes formula (also known as Bayesian filter) by combining predictions with measurements. Together, the sensory and inference compartments form a module that can be stacked to build a hierarchal model using higher-order Markov matrices, enabling the model to learn dynamics that span longer time-scales. This modular implementation allows testing hypotheses related to different learning rules, sensory models and may optionally be used to incorporate actions in the inference process.

Results

We highlight the key features of this model in two simulated experiments.

Effect of Noise In the first experiment we show how dynamics estimates are affected by stochasticity in environmental dynamics as well as "background" noise that could represent noise in the environment or noise introduced by the sensors. In this experiment, we trained six models to learn a 1st order transition matrix (i.e. no hierarchy). All models were trained for a duration of 30 time points. Figure 1 shows on the left-most columns two graphs, one for the true dynamics of a deterministic sequence (top) and a stochastic sequence (bottom). The remaining 3 graphs in each row illustrate the learned dynamics after 300 time points under an ideal sensor (i.e. no noise), low noise levels (Gaussian noise with $\mu = 0, \sigma = 0.3$) and higher noise levels (Gaussian noise with $\mu = 0, \sigma = 0.5$). It should be noted that given the non-parametric nature of our model, sensor noise is not constrained to be of a specific type. In all cases, our approach allowed learning meaningful transitions between states. Furthermore, our model shows that increasing stochasticity in state changes within the environment, resulted in the model relying more on predictions to calculate its belief, whereas clean sensory information led to faster updating of the model dynamics, although not deducible from the figure shown here.

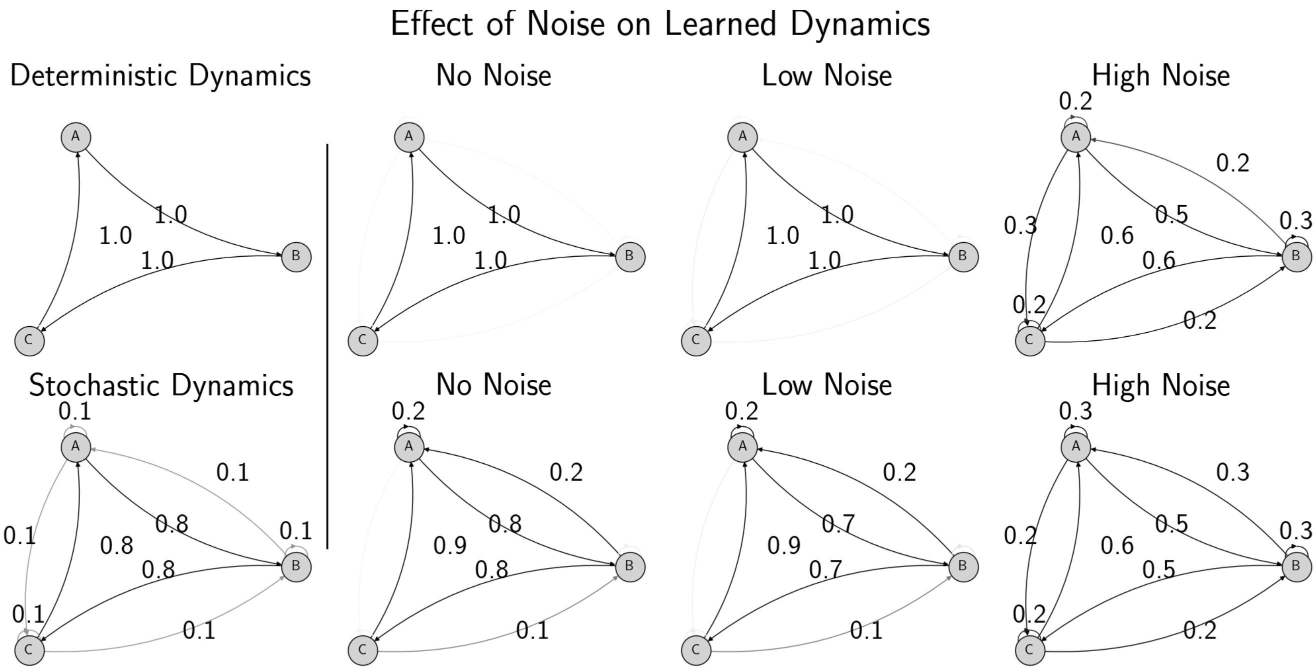

Figure 1: Overview of the true dynamics and the learned internal representations with 3 levels of noise.

Hierarchical Predictions In the second simulation we reproduce prediction error effects of the local global paradigm (see Figure 2) which is commonly employed to investigate the brain's ability to process regularities that span multiple hierarchical levels (Chao, Takaura, Wang, Fujii, & Dehaene, 2018). For this, we built a two-level model where the first level learns 1st order dynamics and the second level learns 3rd order dynamics (i.e. accounting for 3 time-points in the past). We trained the model over 200 time points where the sequence consisted of a sub-sequence of states **AAAB** showing local deviants (acting as global standards) while in the last 50 time-points shown in the figure, the sequence contained 3 local standards / global deviants (**AAAA**). In this model, the higher level received as input not the measurements but the posterior of the lower level. Conversely, the lower level received predictions from the higher level which were combined with the lower level predictions. We show that both local and global (expectation driven) dynamics can be learned using this hierarchical implementation of the model.

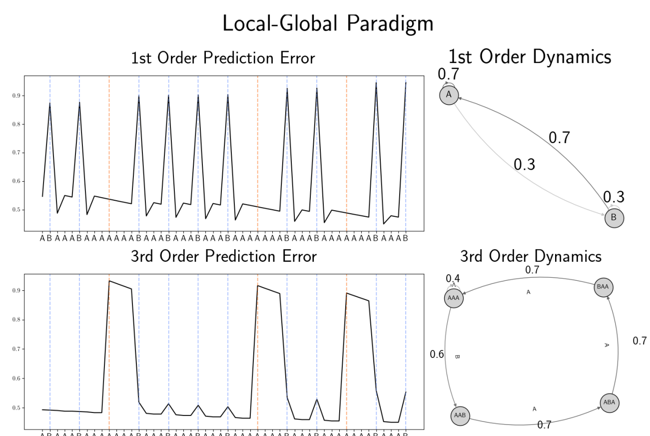

Figure 2: Local and global prediction errors related to the 1st order predictions (top) and 3rd order predictions (bottom). The right graphs show the learned dynamics for the two hierarchical levels in the model.

Discussion

In this work, we showed that exact Bayesian inference and online learning is an appropriate approach for modeling perception in noisy and uncertain environments. Our model hypothesizes the brain to operate on a discrete and restricted state space of the environment and considers sensory characteristics (e.g. receptive fields) and working memory. Our non-parametric model does not make any assumptions about particular sufficient statistics of parametrized distributions and instead models full state distributions using exact online inference. Our model relies only on local Hebbian-like learning rules model updating. This approach of simultaneous learning and inference enables flexible adaptation to the noisy and non-stationary dynamics of the environment. Future work will focus on fitting the model to behavioral data to better understand individual variability during inferential processes, something that could explain individual variability in neuro-psychiatric disorders. To this end, our aim is to provide a user-friendly python package to enable testing of hypotheses related to predictive processing using exact inference and unsupervised learning.

References

Chao, Z. C., Takaura, K., Wang, L., Fujii, N., & Dehaene, S. (2018). Large-scale cortical networks for hierarchical prediction and prediction error in the primate brain. *Neuron*, *100*(5), 1252–1266.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, *11*(2), 127–138.