



Retrospective Inference as a Form of Bounded Rationality, and Its Beneficial Influence on Learning

Thomas H. B. FitzGerald^{1,2,3*}, Will D. Penny^{1,2}, Heidi M. Bonnici¹ and Rick A. Adams^{2,3,4}

¹ School of Psychology, University of East Anglia, Norwich, United Kingdom, ² The Wellcome Trust Centre for Neuroimaging, University College London, London, United Kingdom, ³ Max Planck-UCL Centre for Computational Psychiatry and Ageing Research, London, United Kingdom, ⁴ Department of Computer Science, University College London, London, United Kingdom

OPEN ACCESS

Edited by:

Dimitrije Marković,
Dresden University of
Technology, Germany

Reviewed by:

Maria Chan,
Memorial Sloan Kettering Cancer
Center, United States
Dirk Ostwald,
Freie Universität Berlin, Germany

*Correspondence:

Thomas H. B. FitzGerald
t.fitzgerald@uea.ac.uk

Specialty section:

This article was submitted to
Medicine and Public Health,
a section of the journal
Frontiers in Artificial Intelligence

Received: 12 July 2019

Accepted: 14 January 2020

Published: 18 February 2020

Citation:

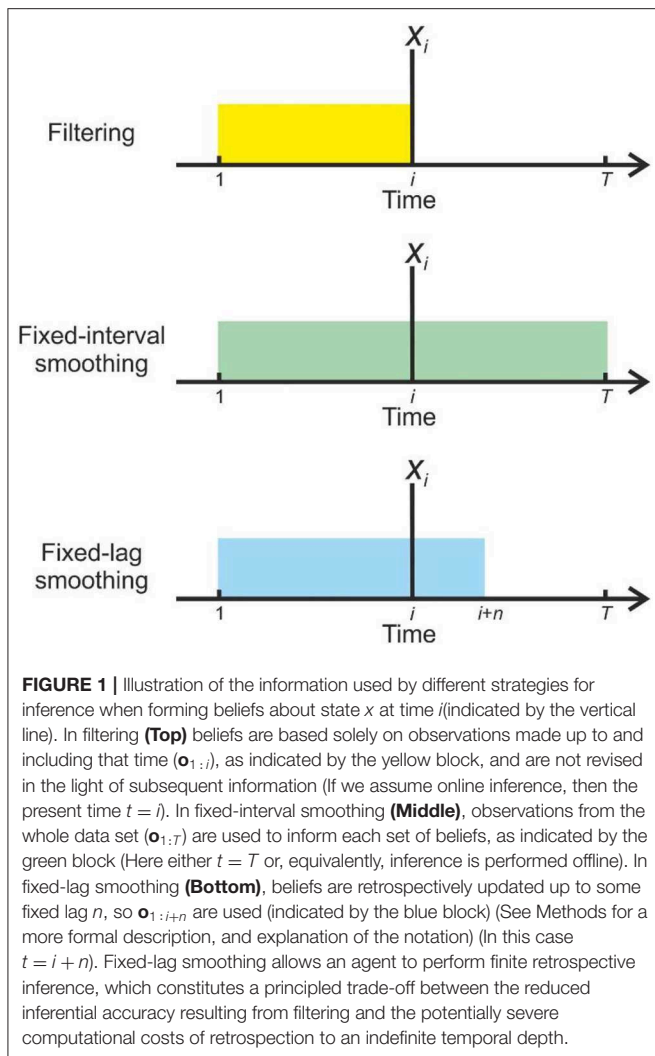
FitzGerald THB, Penny WD,
Bonnici HM and Adams RA (2020)
Retrospective Inference as a Form of
Bounded Rationality, and Its Beneficial
Influence on Learning.
Front. Artif. Intell. 3:2.
doi: 10.3389/frai.2020.00002

Probabilistic models of cognition typically assume that agents make inferences about current states by combining new sensory information with fixed beliefs about the past, an approach known as Bayesian filtering. This is computationally parsimonious, but, in general, leads to suboptimal beliefs about past states, since it ignores the fact that new observations typically contain information about the past as well as the present. This is disadvantageous both because knowledge of past states may be intrinsically valuable, and because it impairs learning about fixed or slowly changing parameters of the environment. For these reasons, in offline data analysis it is usual to infer on every set of states using the entire time series of observations, an approach known as (fixed-interval) Bayesian smoothing. Unfortunately, however, this is impractical for real agents, since it requires the maintenance and updating of beliefs about an ever-growing set of states. We propose an intermediate approach, finite retrospective inference (FRI), in which agents perform update beliefs about a limited number of past states (Formally, this represents online fixed-lag smoothing with a sliding window). This can be seen as a form of bounded rationality in which agents seek to optimize the accuracy of their beliefs subject to computational and other resource costs. We show through simulation that this approach has the capacity to significantly increase the accuracy of both inference and learning, using a simple variational scheme applied to both randomly generated Hidden Markov models (HMMs), and a specific application of the HMM, in the form of the widely used probabilistic reversal task. Our proposal thus constitutes a theoretical contribution to normative accounts of bounded rationality, which makes testable empirical predictions that can be explored in future work.

Keywords: bayesian inference, learning, cognition, retrospective inference, reversal learning, bounded rationality, hidden markov model

INTRODUCTION

To behave adaptively, agents need to continuously update their beliefs about present states of the world using both existing knowledge and incoming sensory information, a process that can be formalized according to the principles of probabilistic inference (von Helmholtz, 1867; Gregory, 1980). This simple insight has generated a large field of inquiry than spans most areas of the mind



and brain sciences and seeks to build probabilistic accounts of cognition (Rao and Ballard, 1999; Friston, 2010; Tenenbaum et al., 2011; Clark, 2012; Pouget et al., 2013; Aitchison and Lengyel, 2016).

In this paper, we take this framework for granted, and consider an important and related problem, that of using new sensory information to update beliefs about the past. This is important because, under conditions of uncertainty, new observations can contain significant information about past states as well as present ones (Corlett et al., 2004; Shimojo, 2014; FitzGerald et al., 2017; Moran et al., 2019).

In offline cognition or data analysis (in which agents are dealing with complete data sets, and are not required to respond to them in real time), it is possible to make inferences about all time points simultaneously (Figure 1).

In other words, one uses every observation to inform every belief about hidden states. This option is unavailable to real, embodied agents because they need to perceive and act in time (online) (Throughout this paper, we will denote the present time with t). They thus need to perform retrospective inference to

increase the accuracy of their beliefs about the past. To perform retrospective inference optimally (or, equivalently in this context, to be strictly rational) it is necessary for an agent to update beliefs about a sequence of states stretching backwards to the beginning of the current task or context, or perhaps even to the beginning of its existence. This sequence is both indefinitely long and constantly growing, and representing and updating these beliefs will thus, in many situations, place intolerable demands on any real organism.

We propose an alternative approach, finite retrospective inference (FRI), in which agents update beliefs about states falling within a limited temporal window stretching into the past (FitzGerald et al., 2017). Selecting the size of this window, and thus the depth of retrospective belief updating constitutes a form of bounded rationality (Simon, 1972; Gigerenzer and Goldstein, 1996; Ortega et al., 2015), since it trades off inferential accuracy against resource costs (e.g., the metabolic and neuronal costs associated with representing beliefs, and the time to perform the calculations). The depth of updating performed by an agent in a particular context might be selected using a form of “metareasoning” in response to environmental demands (Russell and Wefald, 1991; Lieder and Griffiths, 2017). In particular, it is likely that where observations are noisier, and/or temporal dependencies are greater (in other words, where the past remains significant for longer) such strategies will be more advantageous, and are likely to be favored, provided that other constraints allow it. Alternatively, the degree of retrospection might be phenotypically specified (and thus, presumably, selected for during species evolution). In either case, a bounded-rational approach to retrospection has the potential to explain and quantify how humans and other organisms approach but do not attain optimal performance on a number of cognitive tasks.

In addition to its appeal on purely computational grounds, this proposal might help to explain the widespread occurrence of “postdictive” phenomena in perception (Eagleman and Sejnowski, 2000; Shimojo, 2014). A number of such phenomena have been noted, but in all of them perception of an event is influenced by things that only occur afterwards, suggesting a purely retrospective inference on perception (Rao et al., 2001) (Retrospective inference has also been described in the context of associative learning paradigms; Corlett et al., 2004; Moran et al., 2019). There thus seems good reason to believe that a neurobiologically plausible scheme for retrospective inference like FRI may provide valuable insights into real cognitive processes.

FRI differs from existing probabilistic accounts of online cognition (Rao and Ballard, 1999; Ma et al., 2006; Friston and Kiebel, 2009; Glaze et al., 2015; Aitchison and Lengyel, 2016), which typically only consider inferences about present states, an approach known as “Bayesian filtering” (though see Rao et al., 2001; Baker et al., 2017; Friston et al., 2017; Kaplan and Friston, 2018). It thus constitutes a novel hypothesis about cognitive function that extends probabilistic models to subsume a broader range of problems. Importantly, as we will illustrate in the simulations described below, FRI makes testable predictions about behavior and brain activity in real agents that can be tested in future experimental studies.

MATERIALS AND METHODS

Approximating Normative Inference

Consider the situation in which an agent seeks to infer on a series of T time-varying hidden states $\mathbf{x}_{1:T} = \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$ given a set of time-invariant parameters θ that are known with certainty, a series of observations $\mathbf{o}_{1:T} = \{\mathbf{o}_1, \dots, \mathbf{o}_T\}$, and an initial distribution on \mathbf{x}_0 (Both $\mathbf{x}_{1:T}$ and $\mathbf{o}_{1:T}$ are thus random vectors). To simplify our discussion, in what follows we will assume that all the processes under consideration share the following conditional independence properties:

$$p(\mathbf{x}_i | \mathbf{x}_{1:i-1}, \mathbf{o}_{1:T}, \theta) = p(\mathbf{x}_i | \mathbf{x}_{i-1}, \mathbf{o}_{i:T}, \theta), \quad (1)$$

meaning that states at time i depend only upon the immediately preceding states, and are otherwise independent of previous states or observations (this is the Markov property) and that

$$p(\mathbf{o}_i | \mathbf{x}_{1:T}, \mathbf{o}_{1:i-1}, \mathbf{o}_{i+1:T}, \theta) = p(\mathbf{o}_i | \mathbf{x}_i, \theta), \quad (2)$$

meaning that observations depend only on the current states, and not previous states or observations. However, the general principles presented in this paper apply equally in cases where many, if not all, of these properties are relaxed, e.g., in processes with a higher-order temporal structure. We start with a general discussion of retrospective inference, which makes no specification about the nature of states and observations, before discussing a specific instantiation below (the HMM).

By the chain rule of probability, and making use of Equation 1, the joint conditional distribution over all states is given by:

$$p(\mathbf{x}_{1:T} | \mathbf{o}_{1:T}, \mathbf{x}_0, \theta) = \prod_{i=1:T} p(\mathbf{x}_i | \mathbf{o}_{i:T}, \mathbf{x}_{i-1}, \theta). \quad (3)$$

However, inferring on the joint distribution rapidly becomes computationally intractable, and is often unnecessary. Thus, instead of inferring on the joint conditional distribution we can instead infer on the marginal distributions over states at each time point. In other words, infer on the sequence of most likely states rather than the most likely sequence of states. This approach is known as fixed-interval Bayesian smoothing (Sarkka, 2013). The agent can thus be thought of as approximating the joint conditional distribution as:

$$p(\mathbf{x}_{1:T} | \mathbf{o}_{1:T}, \mathbf{x}_0, \theta) \approx \prod_{i=1:T} p(\mathbf{x}_i | \mathbf{o}_{1:T}, \mathbf{x}_0, \theta). \quad (4)$$

This provides a powerful approach for analyzing sequential data, and is widely used in offline data analysis. However, it presents serious practical difficulties for agents performing online inference of the kind that is mandatory for real, embodied agents. That is, where agents have to make inferences, and very likely take actions, whilst the process is unfolding. Specifically, it requires the agent to store and update an ever-growing set of beliefs about the past, resulting in a set of calculations that will rapidly overwhelm the cognitive capacities of plausible embodied agents. This means that “true” rationality, defined here as cognition that accords precisely with the principles of optimal probabilistic

inference, is impossible for real agents, who must instead seek a feasible approximation.

In probabilistic models of online cognition, this is typically achieved by conditioning inference only on past and current observations, an approach known as Bayesian filtering (Sarkka, 2013). This means that agents make inferences of the form:

$$\prod_{i=1:T} p(\mathbf{x}_i | \mathbf{o}_{1:T}, \mathbf{x}_0, \theta) \approx \prod_{i=1:T} p(\mathbf{x}_i | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta). \quad (5)$$

From the perspective of disembodied normative inference, the approximation implied here represents suboptimality. However, for a real cognitive agent, it can be thought of as an unavoidable cost of having to perform inference in time, which necessitates the use of an alternate strategy.

Filtering can be implemented in a straightforward fashion by recursive application of:

$$p(\mathbf{x}_i | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta) = \int \frac{p(\mathbf{o}_i | \mathbf{x}_i, \theta) p(\mathbf{x}_i | \mathbf{x}_{i-1}, \theta)}{p(\mathbf{o}_i | \mathbf{x}_0, \theta)} p(\mathbf{x}_{i-1} | \mathbf{o}_{1:i-1}, \mathbf{x}_0, \theta) d\mathbf{x}_{i-1}. \quad (6)$$

It is thus computationally parsimonious, since it requires only a single set of calculations at each time step and only requires an agent to store fixed beliefs about the past. In the case of first-order processes, this is only about the immediately preceding time step. However, this parsimony comes at a cost, since it reduces the accuracy of an agent’s beliefs about the past, and consequently, as will be discussed later, impairs learning.

To remedy this, an agent that is performing Bayesian filtering, can implement smoothing recursively by performing an additional “backwards pass” through the data:

$$p(\mathbf{x}_i | \mathbf{o}_{1:T}, \mathbf{x}_0, \theta) = p(\mathbf{x}_i | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta) \int \left[\frac{p(\mathbf{x}_{i+1} | \mathbf{x}_i) p(\mathbf{x}_{i+1} | \mathbf{o}_{1:T}, \mathbf{x}_0, \theta)}{p(\mathbf{x}_{i+1} | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta)} \right] d\mathbf{x}_{i+1}. \quad (7)$$

(Use of an integral here and in Equations 8, 10 presupposes that states are continuous-valued. In the case of discrete states, as in the HMM discussed below, this is replaced with a summation). Here $p(\mathbf{x}_i | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta)$ is the state estimate derived from filtering, $p(\mathbf{x}_{i+1} | \mathbf{x}_i)$ is the dynamic model governing transitions between states, $p(\mathbf{x}_{i+1} | \mathbf{o}_{1:T}, \mathbf{x}_0, \theta)$ is the smoothed state estimate at $i + 1$, and $p(\mathbf{x}_{i+1} | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta)$ is the predicted distribution at $i + 1$ given by:

$$p(\mathbf{x}_{i+1} | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta) = \int p(\mathbf{x}_{i+1} | \mathbf{x}_i) p(\mathbf{x}_i | \mathbf{o}_{1:i}, \mathbf{x}_0, \theta) d\mathbf{x}_i. \quad (8)$$

Thus (fixed-interval), smoothing can be carried out in a straightforward manner, beginning with the current state estimate derived from filtering, and working iteratively backwards. Nonetheless, it requires the agent to perform a set of calculations that grows linearly with the time series, and store a similarly growing set of beliefs about past states, and thus introduces significant extra costs for an agent over and above filtering, which are likely to become unsustainable for real agents in ecological contexts. We thus propose that agents make use of

an intermediate strategy, finite retrospective inference, in which they perform retrospective belief updating to a limited degree, in a manner that reflects both the desirability of accurate inference and the need to limit resource (and other) costs.

Finite Retrospective Inference

To implement FRI, we propose that agents perform fixed-lag smoothing, an approach that is intermediate between full (fixed-interval) smoothing and filtering. In fixed-lag smoothing, agents update beliefs about all states within a fixed-length time window that includes the present time but stretches a set distance into the past (**Figure 1**) (FitzGerald et al., 2017). This window moves forward in time at the same rate that observations are gathered, meaning that cognition occurs within a sliding window (In principle the sliding window approach can also be used to infer on the joint distribution of short sequences of states FitzGerald et al., 2017, but we focus on smoothing in this paper for the sake of simplicity). We are unaware of a precedent for this approach in treatments of cognition, however it has been employed in other contexts (Moore, 1973; Cohn et al., 1994; Chen and Tugnait, 2001; Sarkka, 2013). This means that, for a window of length considered at time t , agents approximate the true marginal distribution as follows:

$$\prod_{i=1:T} p(\mathbf{x}_i | \mathbf{o}_{1:T}, \mathbf{x}_0, \boldsymbol{\theta}) \approx \prod_{i=1:T} p(\mathbf{x}_i | \mathbf{o}_{1:i+n-1}, \mathbf{x}_0, \boldsymbol{\theta}). \quad (9)$$

As can be seen by comparing Equations (5, 9), filtering is a special case of fixed-lag smoothing in which $n = 1$. Smoothing can thus be performed by iteratively evaluating.

$$\prod_{i=t-n+1}^t p(\mathbf{x}_i | \mathbf{o}_{1:t}, \mathbf{d}, \boldsymbol{\theta}) = \int p(\mathbf{x}_{t-n} | \mathbf{o}_{1:t-n}, \mathbf{d}, \boldsymbol{\theta}) \prod_{i=t-n+1}^t p(\mathbf{x}_i | \mathbf{o}_{t-n+1:t}, \mathbf{x}_{t-n}, \mathbf{d}, \boldsymbol{\theta}) d\mathbf{x}_{t-n}. \quad (10)$$

This simply requires the agent to track $p(\mathbf{x}_{t-n} | \mathbf{o}_{1:t-n}, \mathbf{d}, \boldsymbol{\theta})$, the filtered estimate of the states that obtain at the timestep immediately preceding the current window. Practically, fixed-lag smoothing can be implemented using Equation (7), with the proviso that backward recursion is only performed $n - 1$ times. In other words, rather than propagating new information right the way back through a time series as is typical in offline applications, it is only propagated to a fixed depth ($n - 1$), limiting the computational cost to the agent. This allows agents to adopt a bounded rational strategy in which they trade off inferential accuracy and computational (and potentially other) costs to select an appropriate depth of processing.

Parameter Learning Using Retrospective Inference

We next consider the more general situation in which there is uncertainty about both states and parameters, and agents must therefore perform learning as well as inference. This is often referred to as a “dual estimation” problem (Wan et al., 1999; Friston et al., 2008; Radillo et al., 2017), and is characteristic

of many real-world situations. To do so, we make the model parameters θ random variables, and condition beliefs about them on a set of fixed hyperparameters λ , such that states and observations are independent of the hyperparameters when conditioned on the parameters, meaning that:

$$p(\mathbf{x}_{1:T}, \boldsymbol{\theta} | \mathbf{o}_{1:T}, \boldsymbol{\lambda}) = p(\mathbf{x}_{1:T} | \mathbf{o}_{1:T}, \boldsymbol{\theta}) p(\boldsymbol{\theta} | \mathbf{o}_{1:T}, \boldsymbol{\lambda}). \quad (11)$$

Learning and inference are inextricably related to one another, since beliefs about states depend on beliefs about parameters, and vice versa. Since parameters are fixed, accurately estimating them involves accumulating evidence across entire time series, and thus beliefs about multiple sets of states. This means that increasing the accuracy of beliefs about the past, through retrospective belief updating, also increases the accuracy of parameter estimation. Crucially, improved parameter estimates will also result in more accurate beliefs about the present and better predictions about the future. Thus, in the context of uncertainty about model parameters, retrospective belief updating is advantageous even for an agent that has no intrinsic interest in the past. This is a very important point, since it argues for the wide importance of retrospective belief updating across a variety of situations and agents.

At a practical level, learning using FRI is very similar to offline learning. We treat each window as a time series in its own right, with λ is replaced by $\tilde{\lambda}$, the sufficient statistics of $p(\boldsymbol{\theta} | \mathbf{o}_{1:t-n}, \lambda)$, which is the posterior distribution over the parameters conditioned on all observations preceding the current window, and perform learning and inference as normal. The use of a sliding window does, however introduce a small additional complexity, since successive windows overlap and thus share data points. Thus, if we treated each window as a separate time series we would count each observation multiple times and, as a result, overweight them. To avoid this, when updating $\tilde{\lambda}$ we only use information about states at the first time-point in the window (i.e., $p(\mathbf{x}_{t-n+1} | \mathbf{o}_{1:t}, \boldsymbol{\theta})$) (A more specific example of this is provided for the HMM below). This also means that only the best available estimate of each set of states (in other words, the estimate that will not be revised in light of future evidence) contributes to stored beliefs about the parameters of the model.

Retrospective Inference in Hidden Markov Models

To illustrate the utility of bounded-rational retrospective inference for an agent, we applied the principles described above to Hidden Markov models (**Figure 2**). In principle though, they apply equally to a broad range of models with alternative properties such as continuous state spaces and higher-order temporal structure. In an HMM, the system moves through a series of T time-varying hidden states, each of which is drawn from a discrete state space of dimension K . Hidden states $\mathbf{x}_{1:T}$ are not observed directly, but instead must be inferred from observed variables. Here we assume that these also discrete, with dimension M , but this need not be the case. Thus, at time t (where $t \in \mathbb{N} : t \in \{1, T\}$), \mathbf{x}_t is a binary vector of length K such that $\sum \mathbf{x}_t = 1$, and similarly \mathbf{o}_t is a binary vector of length M




FIGURE 2 | Bayesian graph illustrating the structure of the Hidden Markov model described in the text (Shaded circles indicate variables with known values, unshaded circles indicate hidden variables). Transitions between hidden states x_0 to x_T are governed by the transition matrix **A**, and are first-order Markovian. Observations o_1 to o_T depend only on the current hidden state and the emission matrix **B**. Where the parameters of **A** and **B** need to be learnt, as depicted here we include appropriate sets of Dirichlet priors, parameterized by the matrices Π^a and Π^b , respectively. Beliefs about the initial hidden state x_0 are governed by the parameter vector **d**.

such that $\sum o_t = 1$ (This can also be described as a multinoulli random variable).

Initial state probabilities are encoded in a row vector **d**, which we will hereafter assume to encode a uniform distribution. Transitions between states are first-order Markovian, and the transition probabilities are encoded in a $K \times K$ matrix **A**, such that:

$$A_{jk} \equiv p(x_{i,k} = 1 | x_{i-1,j} = 1), \tag{12}$$

where, $A_{jk} \in [0, 1]$ and $\sum_k A_{jk} = 1$. This means that each row $A_{j\bullet}$ encodes the transition probabilities from state j to the entire state space. From this it follows (Bishop, 2006) that:

$$p(x_i | x_{i-1}, \mathbf{A}) = \prod_{k=1}^K \prod_{j=1}^K A_{jk}^{x_{i-1,j} x_{i,k}}. \tag{13}$$

Similarly, the $M \times K$ matrix **B** encodes the emission probabilities such that:

$$B_{jk} \equiv p(o_{i,k} = 1 | x_{i,j} = 1), \tag{14}$$

where $B_{jk} \in [0, 1]$ and $\sum_k B_{jk} = 1$. Thus, that each row $B_{j\bullet}$ encodes the probability of each observed variable when in state j , and:

$$p(o_i | x_i, \mathbf{B}) = \prod_{k=1}^M \prod_{j=1}^K B_{jk}^{x_{ij} o_{ik}}. \tag{15}$$

Pure Inference in HMMs

To calculate the smoothed marginal posterior $\gamma(x_i)$ in an HMM, we can make use of the forward-backward algorithm (Rabiner, 1989). This involves recursive forward and backward sweeps, that calculates two quantities $\alpha(x_i)$ and $\beta(x_i)$ for each time point (Bishop, 2006) such that:

$$\begin{aligned} \gamma(x_i) &= p(x_i | o_{1:T}, \mathbf{A}, \mathbf{B}), \\ &= \frac{\alpha(x_i) \beta(x_i)}{\sum \alpha(x_i) \beta(x_i)}, \\ \alpha(x_i) &\equiv p(o_{1:i}, x_i | \mathbf{A}, \mathbf{B}), \\ \beta(x_i) &\equiv p(o_{i+1:T} | x_i, \mathbf{A}, \mathbf{B}). \end{aligned} \tag{16}$$

$\alpha(x_i)$ thus corresponds to the unnormalized filtered posterior, and is given by:

$$\alpha(x_1) = (\mathbf{B}o_1) \circ \mathbf{d}, \tag{17}$$

for the first state, and:

$$\alpha(x_i) = (\mathbf{B}o_i) \circ (\mathbf{A}^T \alpha(x_{i-1})), \tag{18}$$

for all subsequent states. Here \circ denotes the Hadamard or element-wise product. $\beta(x_i)$ is given by:

$$\beta(x_i) = \mathbf{A}(\beta(x_{i+1}) \circ (\mathbf{B}o_{i+1})). \tag{19}$$

To apply the sliding window approach to this model, at each timestep we simply evaluate the filtered posterior using and then perform backward inference a fixed number of steps using. This is the key step that enables the agent to perform FRI by inferring over both the present state and a sequence of previous states stretching a fixed distance into the past.

Dual Estimation in HMMs

To learn the transition probabilities of an HMM we first need to define an additional quantity, the dual-slice marginal $\xi(x_i, x_{i-1})$, which corresponds to the joint probability distribution $p(x_i, x_{i-1} | o_{1:T}, \theta)$ (Baum et al., 1970; Bishop, 2006). It is simple to show that:

$$\xi(x_i, x_{i-1}) \propto \mathbf{A} \circ (\alpha(x_{i-1}) ((\mathbf{B}o_i) \circ \beta(x_i))^T). \tag{20}$$

(For a more detailed exposition of this see Bishop, 2006).

Introducing learning renders exact inference impossible, which necessitates the use of an approximation. Broadly speaking, such approximations fall into two categories: sampling approaches (Andrieu et al., 2003), which are computationally expensive but asymptotically exact, and variational approaches which are more computationally efficient but require the

parameter learning is performed by updating the concentration parameters using the following equations (where t indicates the present time):

$$\begin{aligned}\tilde{\Pi}^a &= \Pi^a + \widehat{\mathbf{M}}_{t-n+1}, \\ \widehat{\Pi}^a &= \tilde{\Pi}^a + \sum_{j=t-n+2}^t \widehat{\mathbf{M}}_j, \\ \tilde{\Pi}^b &= \Pi^b + \widehat{\mathbf{x}}_{t-n+1} \widehat{\mathbf{o}}_{t-n+1}^T, \\ \widehat{\Pi}^b &= \tilde{\Pi}^b + \sum_{j=t-n+2}^t \widehat{\mathbf{x}}_j \widehat{\mathbf{o}}_j^T.\end{aligned}\quad (25)$$

Here $\tilde{\Pi}^a$ and $\tilde{\Pi}^b$ denote the fixed-lag parameters that are incremented across time steps, Π^a and Π^b denote the values of the fixed-lag concentration parameters from the previous time step (in other words, the evidence that has been accumulated prior to the current window), and $\widehat{\Pi}^a$ and $\widehat{\Pi}^b$ denote the full estimates of the concentration parameters based on timesteps 1 to t .

The Probabilistic Reversal Task as a Special Case of the HMM

To illustrate the utility of FRI even for relatively straightforward tasks, we simulated inference and learning on a probabilistic reversal paradigm (Hampton et al., 2006; Glaze et al., 2015; Radillo et al., 2017). Briefly, subjects are required to track an underlying hidden state that occasionally switches between one of two possible values, based on probabilistic feedback (In other words, feedback that is only, for example, 85% reliable). This paradigm is both simple and widely used, and the small state space makes illustrating results in graphical form relatively straightforward. In addition, the fact that the paradigm is widely used makes it an appealing tool for exploring to what extent human subjects actually employ FRI when solving this sort of task. The task can be modeled as an HMM, in which there are only two hidden states, which probabilistically generate one of two possible observations (Hampton et al., 2006; Schlagenhaut et al., 2014; Costa et al., 2015; FitzGerald et al., 2017). The parameter r encodes the probability of a reversal between trials, and ν encodes the reliability of observations. Thus,

$$\begin{aligned}\mathbf{A} &= \begin{bmatrix} 1-r & r \\ r & 1-r \end{bmatrix}, \\ \mathbf{B} &= \begin{bmatrix} \nu & 1-\nu \\ 1-\nu & \nu \end{bmatrix}.\end{aligned}\quad (26)$$

(Introducing learning requires a slight modification of the standard HMM parameter update equations to reflect the symmetry of the \mathbf{A} and \mathbf{B} matrices, as described in the Appendix).

Simulations

Probabilistic Reversal Task

To illustrate the effects of retrospective inference on an agent's beliefs whilst doing the probabilistic reversal task, we simulated

1,000 instantiations of a 256 trial task session, with parameters set as $r = 0.1$ and $\nu = 0.85$, plausible values for real versions of the task (e.g., FitzGerald et al., 2017). For the “pure inference” agent, we set extremely strong (and accurate) prior beliefs about \mathbf{A} and \mathbf{B} by setting initial values of $\Pi^{(a)} = 10^6 \mathbf{A}$ and $\Pi^{(b)} = 10^6 \mathbf{B}$. This has the consequence of effectively fixing these parameters to their prior values (in other words, essentially rendering them fixed parameters). For the “dual estimation” agent, we kept the prior beliefs about \mathbf{B} identical, but set weak priors on the transition matrix of:

$$\Pi^{(a)} = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix}.\quad (27)$$

This has the consequence of allowing agents' beliefs to be determined almost completely by the data they encounter. Window lengths for retrospective inference were set at 1, 2, 4, 8, 16, 32, 64, 128, and 256 trials, and we also simulated an agent performing offline (fixed interval) smoothing for comparison. To assess the accuracy of inference and learning, we calculated the log likelihood assigned to the true sequence of hidden states and the true value of r , calculated using Equation 36 (see Appendix), and averaged these across simulations.

The aim of these simulations is to demonstrate the effects, and potential advantages, of performing FRI for an agent, even on relatively simple tasks. However, establishing whether retrospective inference is in fact a feature of human cognition requires careful experimental validation. This will involve careful model-based analysis of behavioral (and possibly neuroimaging) data collected on appropriate behavioral tasks. We intend to address this in future studies.

Random HMMs

To show that the effects that we illustrate are not due to some specific feature of the probabilistic reversal paradigms, we performed similar simulations, this time using HMMs with three possible hidden states, three possible observations, and randomly generated transition probabilities. We generated 10 such HMMs, and simulated 100 instantiations of each, whilst varying the diagonal terms of the emission matrix \mathbf{B} at intervals of 0.05 between 0.65 and 0.95 (and setting the off-diagonal terms to be equal) (This corresponds to varying the degree of perceptual uncertainty). Prior beliefs for the pure inference and dual estimation agents were set as described for the reversal task, and accuracy was assessed in a similar manner, using Equation 23.

RESULTS

To explore the properties of fixed-lag retrospection in pure inference problems (in other words, ones where no learning is necessary), we simulated behavior on both the probabilistic reversal task and on random HMMs. As expected, in both cases, FRI considerably improved the accuracy of agents' final (offline) beliefs about past hidden states. (Online estimates of current states are identical under all approaches). Strikingly, in both cases, this improvement occurred even when agents only retrospected over short windows (Figure 3), suggesting that, in




FIGURE 3 | Retrospective belief updating improves state estimates during pure inference on a probabilistic reversal task. Top panel: illustration of the first 64 trials of a 256 trial session of the reversal task using different strategies. The final (retrospective) posteriors are shown in blue ($n = 1$, filtering), orange ($n = 2$), and gold ($n = 128$). Black crosses show the true hidden state, and red circles the observations made on each trial. Retrospective belief updating allows agents to infer the true underlying states more accurately. Bottom panel: relative log accuracy of models of different window lengths, averaged across simulated time series (see main text for details) (Accuracy is quantified as the log likelihood assigned to the true sequence of states by the agent, averaged across simulations). This illustrates that, in this context at least, even the use of a very short window leads to significantly more accurate beliefs, but that this benefit saturates relatively rapidly (by about $n = 8$). Thus, for a bounded rational agent performing pure inference, the optimal window length may be surprisingly low, depending on the relevant computational costs.

certain problems at least, a limited capacity for retrospection can yield significantly improves inference.

Simulations of dual estimation problems in which there was uncertainty about r clearly illustrated that retrospective inference increases the accuracy of both retrospective *and* online state estimation, as a result of increased accuracy in parameter learning (Figures 4, 5). One important feature to note is that even when the maximum possible depth of retrospection is employed ($n = 256$), the accuracy of online state estimation always falls significantly short of offline estimation. This indicates the fact that, however great the representational and computational sophistication of an agent, there is always a cost to performing inference online, rather than with a complete data set. If sufficiently high, this cost provides an incentive to perform additional (subsequent) offline processing, perhaps during sleep, and it is conceivable that this might be linked to the extended process of memory consolidation. Similar patterns were observed in the random HMM simulations, supporting the notion that these are general properties of retrospective inference (Figure 6).

DISCUSSION

In this paper, we consider the problem of accurately updating beliefs about the past from the perspective of probabilistic cognition. Specifically, we propose that humans and other agents use finite retrospective inference, in which beliefs about past states are modifiable within a certain temporal window, but are fixed thereafter. We show, using simulations of inference and learning in the context of a probabilistic reversal task, that even a fairly limited degree of retrospection results in significantly improved accuracy of beliefs about both states and parameters. Importantly, the hypothesis that agents perform retrospective inference makes clear predictions about behavior on appropriate tasks that are quantitatively dissociable from those made under the hypothesis that agents use pure filtering. Implementing retrospective inference also makes specific predictions about brain function, since it requires beliefs about past states to be explicitly represented and updated. Our work thus provides testable hypotheses that can be explored in future behavioral and neurobiological studies.


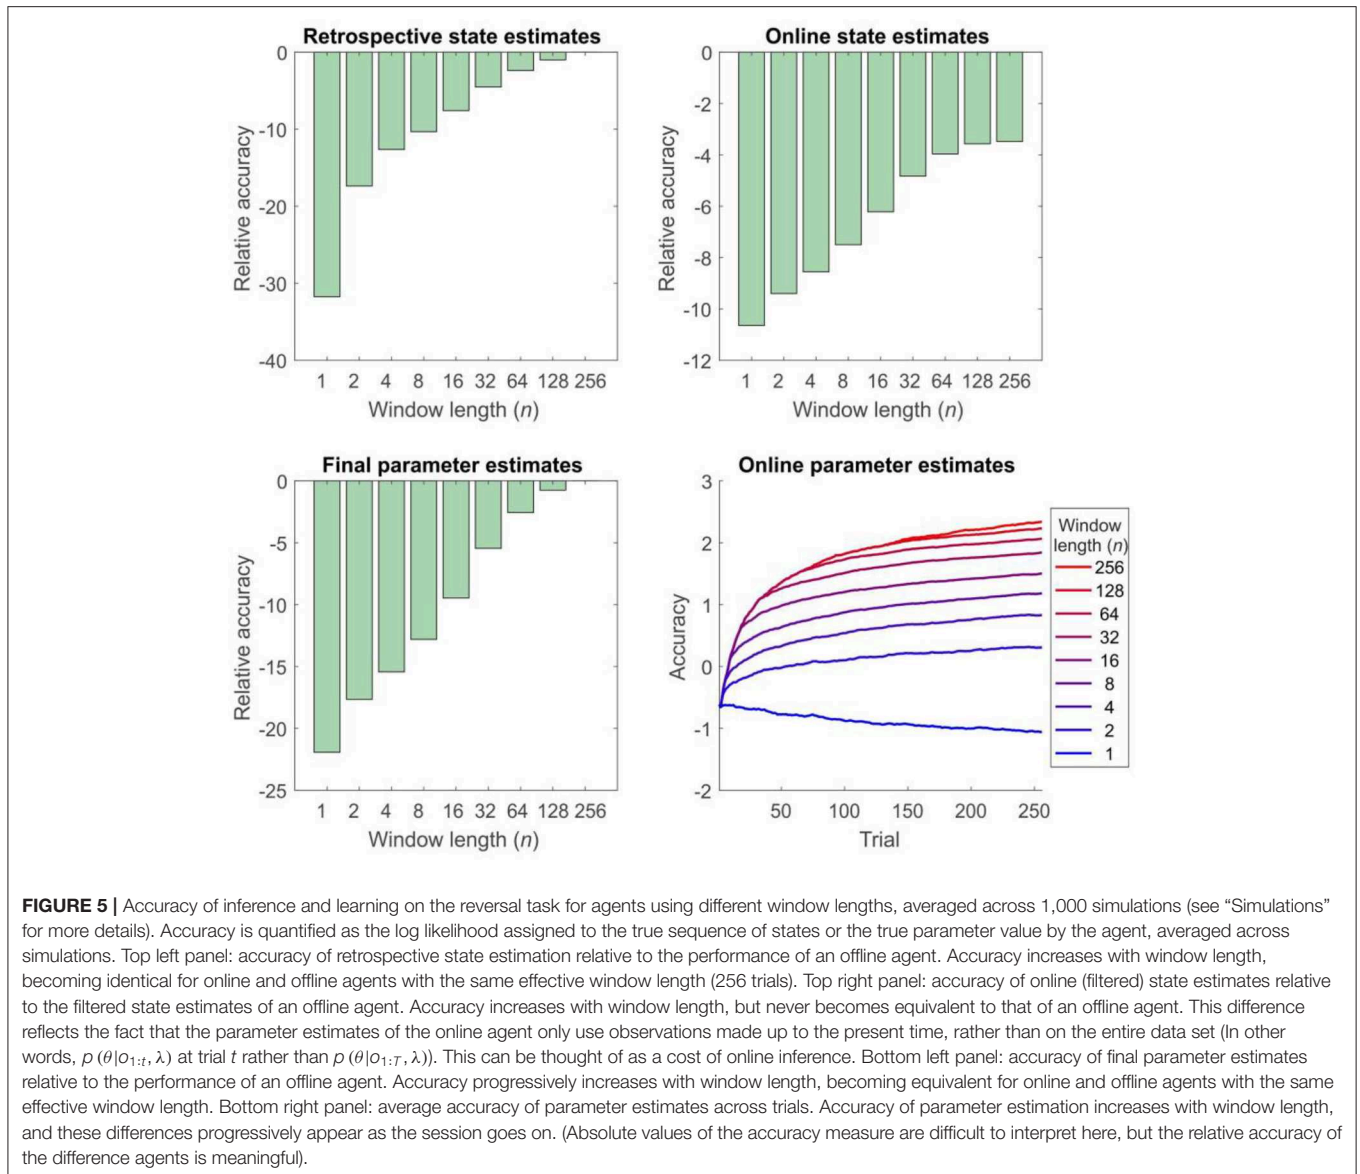


FIGURE 4 | State and parameter estimation for agents performing dual estimation on the first 64 trials of a 256 trial session of the reversal task using different strategies. Top panel: the accuracy of retrospective belief estimates $p(x_t|o_{1:T})$ increases with greater window lengths, but still falls well-short of the performance of an offline agents, who has access to the entire time series simultaneously. Middle panel: the accuracy of online (filtered) beliefs about the current state $p(x_t|o_{1:t})$ subtly but consistently increases with greater window length. Note that this effect is entirely due to the beneficial effects of greater window lengths on parameter learning. Bottom panel: the effect of window length on parameter learning. Estimates of r are derived from Π^a at each timestep (the best estimate available to the agent at that time). With greater window lengths, parameter estimates converge more rapidly on the true value (Estimates from agents performing retrospective inference with windows of length 1, 2, and 128 time steps are shown in blue, orange, and gold, respectively. Estimates from an agent performing offline inference is shown in purple. True hidden states are indicated with black crosses, whilst observations are indicated with red circles. The true value of parameter r is indicated with a dotted black line in 5c).

Perhaps the most significant feature of our simulations is the demonstration that, where there is uncertainty about time-invariant model parameters, finite retrospective inference significantly improves the accuracy of learning. This is important both because these parameters may be of intrinsic interest, and because better learning will result in more accurate beliefs about present and future states. Even if an agent has no intrinsic interest in past events, it still has a clear incentive to perform retrospective inference, since this will allow it to act better in the future. This provides a new twist on the often-advanced hypothesis that the primary function of memory in general, and episodic memory in particular, is to improve predictions about the future (Schacter et al., 2012). Here, in addition to playing a role in constructing imagined future scenarios (Hassabis et al., 2007), the explicit representation of events or episodes in the past may be essential for updating beliefs about

current and future states or learning time-invariant properties of an agent's environment (Baker et al., 2017). In this paper we have not sought to clearly characterize the sorts of problem for which FRI is likely to be most useful. However, this will be extremely important for future work aimed at furnishing empirical evidence for an effect of retrospective inference on learning.

A similar point may be made about the potential importance of retrospective inference for the generation and selection of appropriate cognitive models, a process known as structure learning (Acuña and Schrater, 2010; Braun et al., 2010; Tervo et al., 2016). In this paper, we confine ourselves to considering inference about hidden states and learning about fixed model parameters, but structure learning is an equally important process, and one that is likely to be strongly affected by the depth of retrospective processing employed by an individual. In future

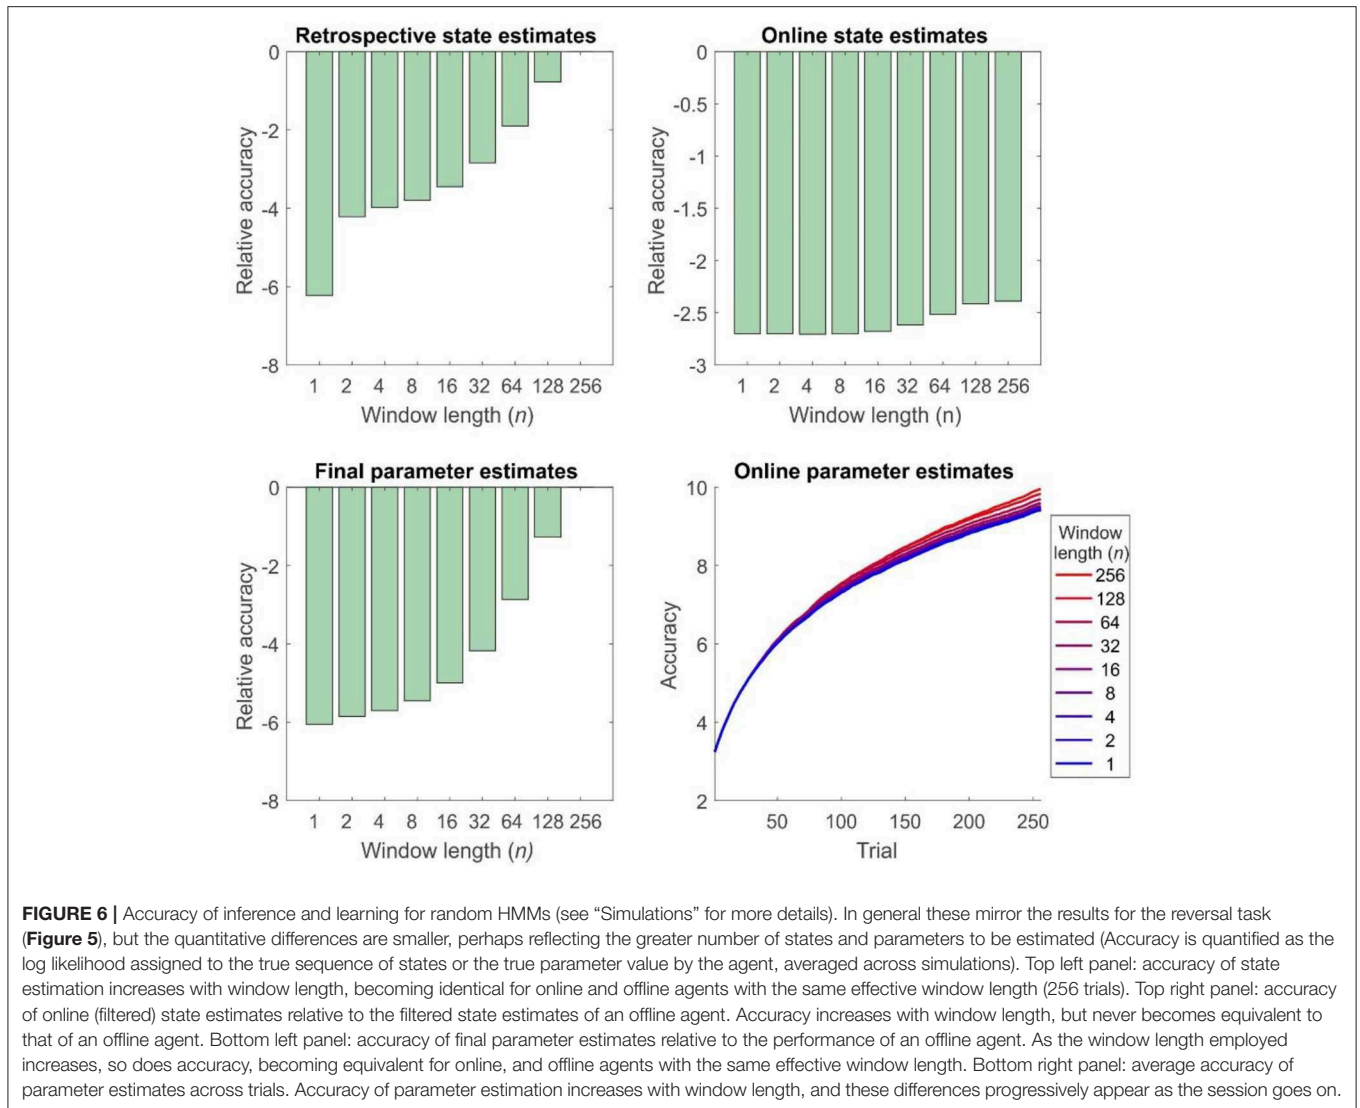


work, we plan to address this explicitly, both through simulations and experimental work.

The specific retrospective inference model we describe here differs importantly from previous approaches to modeling probabilistic reversal tasks (Hampton et al., 2006) and change point detection more generally (Wilson et al., 2010; Radillo et al., 2017) in two key ways, first through the fact that we allow for parameter learning (though see Radillo et al., 2017), and second, because we simulate agents that are able to update beliefs about past states. Both these processes are important for normative behavior, and it will be important to establish how closely human performance across a number of domains reflects this. Retrospective inference has also been considered in the context of reinforcement learning (Moran et al., 2019), and we will explore how to approach similar reward learning problems using our probabilistic framework in future. Similar ideas have

also been explored in the context of active inference and planning (Friston et al., 2017; Kaplan and Friston, 2018), although these have not explored effects on learning.

Our approach differs importantly from models such as the hierarchical Gaussian filter (Mathys et al., 2011), which use higher-level variables operating at longer time scales to provide an implicit time window, but do not make postdictive inferences of the sort discussed here [In fact, retrospective inference has the potential to improve accuracy on tasks involving tracking of higher order variables like volatility (Behrens et al., 2007; Mathys et al., 2011), which is a promising area for future study]. A closer analogy can be drawn with generalized filtering approaches (Friston, 2008), which infer both on the current state and its derivatives (rate of change, acceleration and so on), and require a finite window of data to perform updates. This similarity is something we intend to return to in future work.



Retrospective inference provides a natural explanation for a number of “postdictive” phenomena in perception, in which perception of an event is influenced by other events that only occur afterwards (Eagleman and Sejnowski, 2000; Shimojo, 2014). A classic example of this is the color phi phenomenon (Kolers and von Grünau, 1976). Here, two differently colored dots are briefly displayed to the subject at different spatial locations. If the interval between the flashes is sufficiently short, subjects report perceiving a single moving dot, rather than two separate dots. Critically, they also perceive the color of the dot as changing during motion, meaning that they perceive the second color as occurring before it is presented on screen. This means that information about the color of the second dot has somehow been propagated backwards in (perceptual) time. That such postdictive phenomena might be explained by smoothing has previously been pointed out by Rao et al. (2001), but our proposal builds on this by suggesting a limited window of updating,

as well as highlighting the importance of such belief updating for learning.

The existence of postdictive perceptual phenomena (among other considerations) have led to what is often called the “multiple drafts” account of consciousness (Dennett and Kinsbourne, 1992), in which the contents of conscious are subject to continual revision in the light of new information (at short timescales, at least), and what subjects report is critically dependent upon when they are asked. For example, in the color phi experiment, subjects’ reported perceptual experience would differ if they were asked to report it before the second dot is shown, as opposed to when they are asked to report it afterwards. This accords extremely well with our proposal (at least if we make the further supposition that the contents of consciousness can, in some sense, be identified with the outcome of optimal perceptual inference). Under FRI (unlike filtering), reported perceptual experience will be critically dependent on when the report is made, since online retrospective inference makes beliefs

time-dependent. In other words, my belief about what happened at time t may be different depending on whether you ask me for it at time $t+1$ or time $t+10$. This means that FRI has the potential to provide the computational underpinning of a “probabilistic multiple drafts” model of perceptual experience.

One intriguing possibility raised by FRI is that different individuals might perform retrospective belief updating to different extents, either on particular tasks or in general, and that this might partially explain between-subject differences in performance on particular tasks (see FitzGerald et al., 2017 for evidence of this). Such differences might even help explain facets of psychopathology (Montague et al., 2012). For example, impaired learning due to reduced or absent retrospection might lead to the tendency to form delusional beliefs (Hemsley and Garety, 1986; Corlett et al., 2004; Adams et al., 2013). For example, say someone looked at you in an unusual way—making you feel they were spying on you—but then subsequently ignored you: if you could not use the latter information to revise your initial suspicion, you would be more likely to become paranoid about that person. This idea is supported by the finding of altered neuronal responses in subjects with delusions (as compared with healthy controls) during performance of a retrospective belief updating task (Corlett et al., 2007), and is something we intend to return to in future.

Implementing retrospective inference also has important implications for neurobiology. In particular, since agents need to be able to dynamically update beliefs about past states, they are required to store explicit, ordered representations of the past, and it should be possible to find evidence of this in appropriate neuronal structures (Pezzulo et al., 2014) (For some evidence of this, see Corlett et al., 2004). Intriguingly, this fits extremely well with an extensive literature on hippocampal function (Fortin et al., 2002; Jensen and Lisman, 2005; Pastalkova et al., 2008; Lehn et al., 2009; Penny et al., 2013), a finding supported by the results of our previous study, which found a relation between depth of retrospective processing and gray matter density in the hippocampus (FitzGerald et al., 2017). On the further supposition that retrospective inference is implemented using filtering and smoothing as described above, this leads to the hypothesis that forward and backward sweeps through recently encountered states, as are known to occur in the hippocampus (Diba and Buzsáki, 2007; Pastalkova et al., 2008; Davidson et al., 2009; Wikenheiser and Redish, 2013) may play a key role in retrospective belief updating. What is less clear, at present, is how to implement retrospective inference within established, neurobiologically-grounded accounts of probabilistic inference in the brain (Friston, 2005; Ma et al., 2006; Aitchison and Lengyel,

2016)—though see (Friston et al., 2017) for related suggestions. This is an extremely important question, and one we intend to return to in future work.

Probabilistic models of cognition are an enormously exciting tool for understanding the complex workings of the mind and brain (Clark, 2012; Friston et al., 2013; Pouget et al., 2013; Aitchison and Lengyel, 2016). The ideas we propose represent a development of such approaches to encompass inference about states in the past, as well as the present. On the further hypothesis that the depth of processing employed is flexible and tailored to the demands of a particular problem or environment, such retrospective processing can also be linked to broader notions of bounded rationality (Simon, 1972; Gigerenzer and Goldstein, 1996; Ortega et al., 2015). We show, through simulations of simple environments, that even a limited degree of retrospection can yield significantly more accurate beliefs about both time-varying states and time-invariant parameters, and thus has the potential to support more adaptive, successful behavior to justify its extra resource costs. This makes it a plausible strategy for real, biological agents to employ FRI makes both behavioral and neuronal predictions in a number of contexts and thus naturally suggests further avenues for exploration in future work.

DATA AVAILABILITY STATEMENT

Matlab code for the modeling described is available at <https://github.com/thbfitz/retro-inf>.

AUTHOR CONTRIBUTIONS

TF conceived the study, created the models, and performed the simulations. WP contributed ideas and code. HB and RA contributed ideas. All authors contributed to writing the manuscript.

FUNDING

TF was supported by a European Research Council (ERC) Starting Grant under the Horizon 2020 program (Grant Agreement 804701). This manuscript has been released as a Pre-Print at <https://www.biorxiv.org/content/10.1101/569574v2>.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frai.2020.00002/full#supplementary-material>

REFERENCES

- Acuña, D. E., and Schrater, P. (2010). Structure learning in human sequential decision-making. *PLoS Comput. Biol.* 6:e1001003. doi: 10.1371/journal.pcbi.1001003
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., and Friston, K. J. (2013). The computational anatomy of psychosis. *Front. Psychiatry* 4:47. doi: 10.3389/fpsy.2013.00047
- Aitchison, L., and Lengyel, M. (2016). The hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS Comput. Biol.* 12:e1005186. doi: 10.1371/journal.pcbi.1005186
- Andrieu, C., de Freitas, N., Doucet, A., and Jordan, M. I. (2003). An introduction to MCMC for machine learning. *Mach. Learn.* 50, 5–43. doi: 10.1023/A:1020281327116
- Baker, C. L., Jara-Ettinger, J., Saxe, R., and Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nat. Hum. Behav.* 1:0064. doi: 10.1038/s41562-017-0064
- Baum, L. E., Petrie, T., Soules, G., and Weiss, N. (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of

- markov chains. *Ann. Math. Stat.* 41, 164–171. doi: 10.1214/aoms/1177697196
- Beal, M. J. (2003). *Variational Algorithms for Approximate Bayesian Inference* (Ph.D. thesis). University College London, London, United Kingdom.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221. doi: 10.1038/nn1954
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York, NY: Springer.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: a review for statisticians. *J. Am. Stat. Assoc.* 112, 859–877. doi: 10.1080/01621459.2017.1285773
- Braun, D. A., Mehring, C., and Wolpert, D. D. M. (2010). Structure learning in action. *Behav. Brain Res.* 206, 157–165. doi: 10.1016/j.bbr.2009.08.031
- Chen, B., and Tugnat, J. K. (2001). Tracking of multiple maneuvering targets in clutter using IMM/JPDA filtering and fixed-lag smoothing. *Automatica* 37, 239–249. doi: 10.1016/S0005-1098(00)00158-8
- Clark, A. (2012). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* 36, 181–204. doi: 10.1017/S0140525X12000477
- Cohn, S. E., Sivakumaran, N. S., Todling, R., Cohn, S. E., Sivakumaran, N. S., and Todling, R. (1994). A fixed-lag kalman smoother for retrospective data assimilation. *Mon. Weather Rev.* 122, 2838–2867. doi: 10.1175/1520-0493(1994)122<2838:AFLKSF>;2.0.CO;2
- Corlett, P. R., Aitken, M. R., Dickinson, A., Shanks, D. R., Honey, G. D., Honey, R. A., et al. (2004). Prediction error during retrospective reevaluation of causal associations in humans: fMRI evidence in favor of an associative model of learning. *Neuron* 44, 877–888. doi: 10.1016/S0896-6273(04)00756-1
- Corlett, P. R., Murray, G. K., Honey, G. D., Aitken, M. R., Shanks, D. R., Robbins, W., et al. (2007). Disrupted prediction-error signal in psychosis: evidence for an associative account of delusions. *Brain* 130, 2387–2400. doi: 10.1093/brain/awm173
- Costa, V. D., Tran, V. L., Turchi, J., and Averbeck, B. B. (2015). Reversal learning and dopamine: a bayesian perspective. *J. Neurosci.* 35, 2407–2416. doi: 10.1523/JNEUROSCI.1989-14.2015
- Davidson, T. J., Kloosterman, F., and Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron* 63, 497–507. doi: 10.1016/j.neuron.2009.07.027
- Dennett, D. C., and Kinsbourne, M. (1992). Time and the observer: the where and when of consciousness in the brain. *Behav. Brain Sci.* 15, 183–201. doi: 10.1017/S0140525X00068229
- Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* 10, 1241–1242. doi: 10.1038/nn1961
- Eagleman, D. M., and Sejnowski, T. J. (2000). Motion integration and postdiction in visual awareness. *Science* 287, 2036–2038. doi: 10.1126/science.287.5460.2036
- FitzGerald, T. H., Hämmerer, D., Friston, K. J., Li, S.-C., and Dolan, R. J. (2017). Sequential inference as a mode of cognition and its correlates in frontoparietal and hippocampal brain regions. *PLoS Comput. Biol.* 13:e1005418. doi: 10.1371/journal.pcbi.1005418
- Fortin, N. J., Agster, K. L., and Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nat. Neurosci.* 5, 458–462. doi: 10.1038/nn834
- Friston, K. (2005). A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 360, 815–836. doi: 10.1098/rstb.2005.1622
- Friston, K. (2008). Variational filtering. *Neuroimage* 41, 747–766. doi: 10.1016/j.neuroimage.2008.03.017
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017). Active inference: a process theory. *Neural Comput.* 29, 1–49. doi: 10.1162/NECO_a_00912
- Friston, K., and Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 364, 1211–1221. doi: 10.1098/rstb.2008.0300
- Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., and Dolan, R. J. (2013). The anatomy of choice: active inference and agency. *Front. Hum. Neurosci.* 7:598. doi: 10.3389/fnhum.2013.00598
- Friston, K., Trujillo-Barreto, N., and Daunizeau, J. (2008). DEM: a variational treatment of dynamic systems. *Neuroimage* 41, 849–885. doi: 10.1016/j.neuroimage.2008.02.054
- Friston, K. J. (2010). The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138. doi: 10.1038/nrn2787
- Gigerenzer, G., and Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychol. Rev.* 103, 650–669. doi: 10.1037/0033-295X.103.4.650
- Glaze, C. M., Kable, J. W., and Gold, J. I. (2015). Normative evidence accumulation in unpredictable environments. *Elife* 4:e08825. doi: 10.7554/eLife.08825.019
- Gregory, R. L. (1980). Perceptions as hypotheses. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 290, 181–197. doi: 10.1098/rstb.1980.0090
- Hampton, A. N., Bossaerts, P., and O’Doherty, J. P. (2006). The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *J. Neurosci.* 26, 8360–8367. doi: 10.1523/JNEUROSCI.1010-06.2006
- Hassabis, D., Kumaran, D., and Maguire, E. A. (2007). Using imagination to understand the neural basis of episodic memory. *J. Neurosci.* 27, 14365–14374. doi: 10.1523/JNEUROSCI.4549-07.2007
- Hemsley, D. R., and Garety, P. A. (1986). The formation of maintenance of delusions: a bayesian analysis. *Br. J. Psychiatry* 149, 51–56. doi: 10.1192/bjp.149.1.51
- Jensen, O., and Lisman, J. E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends Neurosci.* 28, 67–72. doi: 10.1016/j.tins.2004.12.001
- Kaplan, R., and Friston, K. J. (2018). Planning and navigation as active inference. *Biol. Cybern.* 112, 323–343. doi: 10.1007/s00422-018-0753-2
- Kolers, P. A., and von Grünau, M. (1976). Shape and color in apparent motion. *Vision Res.* 16, 329–335. doi: 10.1016/0042-6989(76)90192-9
- Lehn, H., Steffenach, H.-A., van Strien, N. M., Veltman, D. J., Witter, M. P., and Häberg, A. K. (2009). A specific role of the human hippocampus in recall of temporal sequences. *J. Neurosci.* 29, 3475–3484. doi: 10.1523/JNEUROSCI.5370-08.2009
- Lieder, F., and Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychol. Rev.* 124, 762–794. doi: 10.1037/rev0000075
- Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438. doi: 10.1038/nn1790
- Mackay, D. J. C. (1997). *Ensemble Learning for Hidden Markov Models*. Technical report, Cavendish Laboratory, University of Cambridge.
- Mathys, C., Daunizeau, J., Friston, K. J., and Stephan, K. E. (2011). A bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* 5:39. doi: 10.3389/fnhum.2011.00039
- Montague, P. R., Dolan, R. J., Friston, K. J., and Dayan, P. (2012). Computational psychiatry. *Trends Cogn. Sci.* 16, 72–80. doi: 10.1016/j.tics.2011.11.018
- Moore, J. B. (1973). Discrete-time fixed-lag smoothing algorithms. *Automatica* 9, 163–173. doi: 10.1016/0005-1098(73)90071-X
- Moran, R., Keramati, M., Dayan, P., and Dolan, R. J. (2019). Retrospective model-based inference guides model-free credit assignment. *Nat. Commun.* 10:750. doi: 10.1038/s41467-019-08662-8
- Ortega, P. A., Braun, D. A., Dyer, J., Kim, K.-E., and Tishby, N. (2015). Information-theoretic bounded rationality. *arXiv* 1512.06789.
- Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science* 321, 1322–1327. doi: 10.1126/science.1159775
- Penny, W. D., Zeidman, P., and Burgess, N. (2013). Forward and backward inference in spatial cognition. *PLoS Comput. Biol.* 9:e1003383. doi: 10.1371/journal.pcbi.1003383
- Pezzulo, G., van der Meer, M. A., Lansink, C. S., and Pennartz, C. M. (2014). Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn. Sci.* 18, 647–657. doi: 10.1016/j.tics.2014.06.011
- Pouget, A., Beck, J. M., Ma, W. J., and Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16, 1170–1178. doi: 10.1038/nn.3495
- Rabiner, L. R. (1989). A tutorial on hidden markov models and selected applications in speech recognition. *Proc. IEEE* 77, 257–286. doi: 10.1109/5.18626

- Radillo, A. E., Veliz-Cuba, A., Josić, K., and Kilpatrick, Z. P. (2017). Evidence accumulation and change rate inference in dynamic environments. *Neural Comput.* 29, 1561–1610. doi: 10.1162/NECO_a_00957
- Rao, R. P., and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87. doi: 10.1038/4580
- Rao, R. P. N., Eagleman, D. M., and Sejnowski, T. J. (2001). Optimal smoothing in visual motion perception. *Neural Comput.* 13, 1243–1253. doi: 10.1162/08997660152002843
- Russell, S., and Wefald, E. (1991). Principles of metareasoning. *Artif. Intell.* 49, 361–395. doi: 10.1016/0004-3702(91)90015-C
- Sarkka, S. (2013). *Bayesian Filtering and Smoothing*. Cambridge: Cambridge University Press.
- Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., and Szpunar, K. K. (2012). The future of memory: remembering, imagining, and the brain. *Neuron* 76, 677–694. doi: 10.1016/j.neuron.2012.11.001
- Schlagenhauf, F., Huys, Q. J., Deserno, L., Rapp, M. A., Beck, A., Heinze, H. J., et al. (2014). Striatal dysfunction during reversal learning in unmedicated schizophrenia patients. *Neuroimage* 89, 171–180. doi: 10.1016/j.neuroimage.2013.11.034
- Shimojo, S. (2014). Postdiction: its implications on visual awareness, hindsight, and sense of agency. *Front. Psychol.* 5:196. doi: 10.3389/fpsyg.2014.00196
- Simon, H. A. (1972). “Theories of bounded rationality,” in *Decision and Organization*, eds C. McGuire and R. Radner (Amsterdam: North-Holland), 161–176.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science* 331, 1279–1285. doi: 10.1126/science.1192788
- Tervo, D. G. R., Tenenbaum, J. B., and Gershman, S. J. (2016). Toward the neural implementation of structure learning. *Curr. Opin. Neurobiol.* 37, 99–105. doi: 10.1016/j.conb.2016.01.014
- von Helmholtz, H. (1867). *Handbuch der Physiologischen Optik*. Leipzig: Leopold Voss.
- Wan, E. A., Van Der Merwe, R., and Nelson, A. T. (1999). “Dual estimation and the unscented transformation,” in *Proceedings of the 12th International Conference on Neural Information Processing Systems* (Cambridge, MA: MIT Press), 666–672.
- Wikenheiser, A. M., and Redish, A. D. (2013). The balance of forward and backward hippocampal sequences shifts across behavioral states. *Hippocampus* 23, 22–29. doi: 10.1002/hipo.22049
- Wilson, R. C., Nassar, M. R., and Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Comput.* 22, 2452–2476. doi: 10.1162/NECO_a_00007

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 FitzGerald, Penny, Bonnici and Adams. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.