

Multi-Modal and Multi-Factor Branching Time Active Inference.

Abstract

Active inference is a state-of-the-art framework for modelling the brain that explains a wide range of mechanisms. Recently, two versions of branching time active inference (BTAI) have been developed to handle the exponential (space and time) complexity class that occurs when computing the prior over all possible policies up to the time horizon. However, those two versions of BTAI still suffer from an exponential complexity class w.r.t the number of observed and latent variables being modelled. We resolve this limitation by allowing each observation to have its own likelihood mapping, and each latent variable to have its own transition mapping. The implicit mean field approximation was tested — in terms of its efficiency and computational cost — using a dSprites environment in which the metadata of the dSprites dataset was used as input to the model. In this setting, earlier implementations of branching time active inference (namely, $BTAI_{VMP}$ and $BTAI_{BF}$) underperformed in relation to the mean field approximation ($BTAI_{3MF}$); both in terms of performance and computational efficiency. Specifically, $BTAI_{VMP}$ was able to solve 96.9% of the task in 5.1 seconds, and $BTAI_{BF}$ was able to solve 98.6% of the task in 17.5 seconds. Our new approach outperformed both of its predecessors by solving the task completely (100%) in only 2.559 seconds.

Keywords: Branching Time Active Inference, Monte-Carlo Tree Search, Belief Propagation, Bayesian Prediction, Temporal Slice

1. Introduction

Active inference applies the free energy principle to generative models with actions (Friston et al., 2016; Costa et al., 2020; Champion et al., 2021b) and can be regarded as a form of planning as inference (Botvinick and Toussaint, 2012). This framework has successfully explained a wide range of neuro-cognitive phenomena, such as habit formation (Friston et al., 2016), Bayesian surprise (Itti and Baldi, 2009), curiosity (Schwartenbeck et al., 2018), and dopaminergic discharges (FitzGerald et al., 2015). It has also been applied to a variety of tasks, such as animal navigation (Fountas et al., 2020), robotic control (Pezato et al., 2020; Sancaktar et al., 2020), the mountain car problem (Çatal et al., 2020), the game DOOM (Cullen et al., 2018) and the cart pole problem (Millidge, 2019).

However, active inference suffers from an exponential (space and time) complexity class that occurs when computing the prior — i.e., expected free energy — over all possible policies up to the time horizon. Recently, two versions of branching time active inference (BTAI) based on Monte-Carlo tree search (Browne et al., 2012) have been developed to handle this exponential growth. In the original formulation of the framework (Champion et al., 2022b,a), inference was performed using the variational message passing (VMP) algorithm (Winn and Bishop, 2005; Champion et al., 2021b). In a follow up paper, VMP was then replaced by a Bayesian filtering (Fox et al., 2003) scheme leading to a faster inference process (Champion et al., 2021a).

In this paper, we develop an extension of Branching Time Active Inference (BTAI), to allow modelling of several modalities as well as several latent states. Indeed, even if the Bayesian filtering version of Branching Time Active Inference ($BTAI_{BF}$) is fast, its modelling capacity

is limited to one observation and one hidden state. Consequently, if one wanted to model n latent states S_t^1, \dots, S_t^n , then those n latent states would have to be encoded into one latent state X representing all possible configurations of the n latent states S_t^1, \dots, S_t^n . Unfortunately, the total number of configurations is given by:

$$\#X = \prod_{i=1}^n \#S_t^i \geq 2^n,$$

where $\#X$ is the number of possible values taken by X , and similarly $\#S_t^i$ is the number of possible values taken by S_t^i . The above inequality is obtained by realizing that $\#S_t^i \geq 2$, and is problematic in practice because $\#X$ grows exponentially with the number of latent states n being modelled. Also, note that in practice this exponential growth may be way worse than 2^n . For example, if one were to model the five modalities of the dSprites environment (c.f. Section 3.1), the total number of configurations would be:

$$\begin{aligned} \#S_t^y \times \#S_t^x \times \#S_t^{scale} \times \#S_t^{orientation} \times \#S_t^{scale} \\ = 33 \times 32 \times 3 \times 40 \times 6 = 760,320 \gg 2^5 = 32. \end{aligned}$$

A similar exponential explosion also appears when trying to model several modalities O_t^1, \dots, O_t^m using a single one Y , i.e.

$$\#Y = \prod_{i=1}^m \#O_t^i \geq 2^m,$$

where $\#Y$ is the number of possible values taken by Y , and similarly $\#O_t^i$ is the number of possible values taken by O_t^i . Note, throughout this paper, we will use the term *states* to refer to the latent states of the model at a specific

time step, e.g., S_t^1, \dots, S_t^n for time step t . Additionally, we will use the terms *state configurations* or *values* to refer to particular values taken by the latent variables.

Note, the description of the ensuing implementation of planning under active inference is framed from a machine learning perspective; with a focus on computational efficiency afforded by certain mean field approximations. We therefore take care to compare the efficiency of the new scheme with respect to benchmarks established by previous formulations. However, the implications of this scheme reach beyond machine learning and speak to the kinds of factorizations that may be employed by the brain. We will foreground these implementations in the discussion; in terms of functional specialization and implications for functional brain architectures.

The present paper finesses these exponential problems by applying a judicious mean field approximation to future states and observations. This allows the modelling of several observations and latent states, while providing an easy to use framework based on a high-level notation, which allows the user to create models by simply declaring the variables it contains, and the dependencies between those variables. Then, the framework performs the inference process automatically. Appendix A shows an example of how to implement a custom $BTAI_{3MF}$ agent using our framework. In section 2, we describe the theory underlying our approach. Importantly, $BTAI_{3MF}$ takes advantage of the generative model structure to perform inference efficiently using a mixture of belief propagation (Yedidia, 2011; Friston et al., 2017; Kschischang et al., 2001) and forward predictions as will be explained in Section 2.3. The name $BTAI_{3MF}$ is an abbreviation for $BTAI_{MMMMF}$ that stands for: Multi-Modal and Multi-Factor Branching Time Active Inference. Next, in Section 2.4, we provide the definition of the expected free energy in the context of our new approach, and in Section 2.5, we describe the planning algorithm used to expand the generative model dynamically. Then, in Section 3, we compare $BTAI_{3MF}$ to $BTAI_{VMP}$ and $BTAI_{BF}$, and demonstrate empirically that $BTAI_{3MF}$ outperformed both $BTAI_{VMP}$ and $BTAI_{BF}$ on the dSprites environment, which requires the modelling of many latent states and modalities. Finally, Section 4 concludes this paper by summarizing our approach and results.

2. Theory of $BTAI_{3MF}$

In this section, we introduce the mathematical foundation of $BTAI_{3MF}$. To simplify the graphical representation of our generative model, we first introduce a notion of ‘temporal slice’. Then, we build on this idea to describe the

generative model of $BTAI_{3MF}$. Next, we explain how belief updates are performed using a mixture of belief propagation and forward predictions. Afterwards, we provide the definition of the expected free energy for this new generative model. Finally, we describe the planning algorithm used to dynamically expand the generative model, and the action selection process.

2.1 Temporal slice

A temporal slice $TS_J = \{O_J^1, \dots, O_J^{\#O}, S_J^1, \dots, S_J^{\#S}\}$ is a set of random variables indexed by a sequence of actions J . Each random variable of the temporal slice represents either an observation O_J^o or a latent state S_J^s . The index of the temporal slice corresponds to the sequence of actions that lead to this temporal slice. By definition, if J is an empty sequence, i.e., $J = \emptyset$, then TS_J is the temporal slice of the present time step t , also denoted TS_t . Within a temporal slice TS_J , an observation O_J^o depends on a number of latent states $\rho_J^o \subseteq \{S_J^s \mid s = 1, \dots, \#S\}$, such that $P(O_J^o \mid \rho_J^o)$ is a factor in the generative model. Given an action \mathbf{a} and a sequence of actions J , we let $I = J::\mathbf{a}$ be the sequence of actions obtained by appending the action \mathbf{a} at the end of the sequence of actions J . If $I = J::\mathbf{a}$, then the temporal slice TS_J can be the parent of TS_I . This means that a latent state S_I^s in TS_I can depend on the latent states $\rho_I^s \subseteq \{S_J^s \mid s = 1, \dots, \#S\}$ in TS_J , such that $P(S_I^s \mid \rho_I^s)$ is a factor in the generative model. The concept of temporal slice is illustrated in Figure 1, and Figure 2 depicts a more compact representation of the content of Figure 1.

In effect, the temporal slice formulation combines latent states and their consequent observations into a single random variable. Note that this implicitly treats observations as (latent) random variables. This is a key aspect of active inference; in the sense that observable outcomes in the future have not yet been observed and therefore become random variables. This means that the free energy functionals — used to quantify priors over policies — require an expectation over both (future) hidden states and observations. This is why the requisite function is called ‘expected’ free energy.




Figure 1: This figure illustrates two temporal slices TS_t and TS_I , which are depicted by rectangles with thick border. Within each temporal slice, plate notation is used to generate $\#S$ latent states and $\#O$ observations. The dashed lines that connect two random variables from two different plates are new to this paper, and represent an arbitrary connectivity between the two sets of random variables generated by the plates. For example, the dashed line from S_t^s to O_t^o , means that for each observation O_t^o , the parents of O_t^o denoted ρ_t^o is a subset of $\{S_t^s \mid s = 1, \dots, \#S\}$, i.e., the generative model contains the factor $P(O_t^o \mid \rho_t^o)$ where $\rho_t^o \subseteq \{S_t^s \mid s = 1, \dots, \#S\}$.




Figure 2: This figure illustrates the two temporal slices TS_t and TS_I from Figure 1 in a more compact fashion. Since O_t^o is an observed variable for all $o \in \{1, \dots, \#O\}$, the square representing TS_t has a gray background. In contrast, the square representing TS_I has a white background because O_I^o is a latent variable for all $o \in \{1, \dots, \#O\}$.

2.2 Generative model

In this section, we build upon the notion of temporal slice to describe the full generative model. Intuitively, the probability of the entire generative model is the product of the probability of each temporal slice within the model. This includes the current temporal slice TS_t and the future temporal slices TS_I for all $I \in \mathbb{I}$, where \mathbb{I} is the set of all multi-indices expanded during the tree search (c.f., Section 2.5). Within each temporal slice, there are $\#O$ observations and $\#S$ latent states. Each observation depends on a subset of the latent states. Moreover, each latent state depends on a subset of the latent states

of the parent temporal slice. Note, the current temporal slice TS_t does not have any parents, therefore its latent state does not depend on anything. In other words, the model makes the Markov assumption, i.e., each state only depends on the states at the previous time step. More formally, the generative model is defined as:

$$\begin{aligned}
 P(O_t, S_t, O_{\mathbb{I}}, S_{\mathbb{I}}) &= P(TS_t) \prod_{I \in \mathbb{I}} P(TS_I) \\
 &= \underbrace{\prod_{o=1}^{\#O} P(O_t^o \mid \rho_t^o) \prod_{s=1}^{\#S} P(S_t^s)}_{\text{current temporal slice } TS_t} \\
 &\quad \prod_{I \in \mathbb{I}} \left[\underbrace{\prod_{o=1}^{\#O} P(O_I^o \mid \rho_I^o) \prod_{s=1}^{\#S} P(S_I^s \mid \rho_I^s)}_{\text{future temporal slice } TS_I} \right]
 \end{aligned}$$

where t is the current time step, ρ_t^x is the set of parents of X_t^x , $O_t = \{O_t^o \mid o = 1, \dots, \#O\}$ is the set of all observations at time t , $O_I = \{O_I^o \mid o = 1, \dots, \#O\}$ is the set of all future observations that would be observed after performing the sequence of actions I , $O_{\mathbb{I}} = \cup_{I \in \mathbb{I}} O_I$ is the set of all future observations contained in the temporal slices expanded during the tree search (c.f., Section 2.5), $S_t = \{S_t^s \mid s = 1, \dots, \#S\}$ is the set of all latent states at time t , $S_I = \{S_I^s \mid s = 1, \dots, \#S\}$ is the set of random variables describing the future latent states after performing the sequence of actions I , $S_{\mathbb{I}} = \cup_{I \in \mathbb{I}} S_I$ is the set of latent variables representing all future states contained in the temporal slices expanded during the tree search (c.f., Section 2.5). Importantly, the above generative model has to satisfy:

- $\forall I \in \mathbb{I}, \forall o \in \{1, \dots, \#O\}, \rho_I^o \subseteq S_I$;
- $\forall I::\mathbf{a} \in \mathbb{I}, \forall s \in \{1, \dots, \#S\}, \rho_I^{\mathbf{a}:s} \subseteq S_I$, also, if $I = \emptyset$ then by definition $S_I \triangleq S_t$.

Additionally, we define the factors of the generative model as:

$$\begin{aligned}
 P(O_t^o \mid \rho_t^o) &= \text{Cat}(\mathbf{A}^o), & P(S_t^s) &= \text{Cat}(\mathbf{D}_t^s), \\
 P(O_I^o \mid \rho_I^o) &= \text{Cat}(\mathbf{A}^o), & P(S_I^s \mid \rho_I^s) &= \text{Cat}(\mathbf{B}_I^s),
 \end{aligned}$$

where \mathbf{A}^o is the tensor modelling the likelihood mapping of the o -th observation, \mathbf{D}_t^s is the vector modelling the prior over the s -th latent state at time t (see below for details), \mathbf{B}_I^s is the tensor modelling the transition mapping of the s -th latent state under each possible action, \mathbf{B}_I^s is the tensor modelling the transition mapping of the s -th latent state under the last action I_{last} of the sequence I , i.e., $\mathbf{B}_I^s = \mathbf{B}^s(\cdot, \dots, \cdot, I_{\text{last}})$. Also, note that at the beginning of a trial, i.e., when $t = 0$, \mathbf{D}_t^s is a vector that encodes the modeller's understanding of the

task. Afterwards, when $t > 0$, \mathbf{D}_t^s is a vector containing the parameters of the posterior over hidden states according to the observations made and actions taken so far, i.e., $P(S_t^s) \triangleq P(S_t^s | O_{0:t-1}, A_{0:t-1}) = \text{Cat}(\mathbf{D}_t^s)$ for all $s \in \{1, \dots, \#S\}$. Finally, Figure 3 illustrates the full generative model using the notion of temporal slices.




Figure 3: This figure illustrates the full generative model of $BTAI_{3MF}$. The temporal slices depicted in light gray correspond to temporal slices that have not yet been explored by the planning algorithm, c.f., Section 2.5. The numbers between parentheses correspond to the sequence of actions performed to reach the temporal slice.

2.3 Belief updates: the inference and prediction (IP) algorithm

The IP algorithm is composed of two steps, i.e., the inference step (or I-step) and the prediction step (or P-step). The goal of the I-step is to compute the posterior beliefs over all the latent variables at time t . In other words, the goal of the I-step is to compute: $P(S_t^s | O_t), \forall s \in \{1, \dots, \#S\}$. The P-step takes as inputs the posterior beliefs over all the latent variables corresponding to the states of the system after performing a sequence of actions I , and an action \mathbf{a} to be performed next. The goal of the P-step is to compute the posterior beliefs over all the latent variables corresponding to the future states and observations after performing the sequence of actions $I::\mathbf{a}$, where $I::\mathbf{a}$ is the sequence of actions obtained by adding the action \mathbf{a} at the end of the sequence of actions I . In other words, given $P(S_I^s | O_t), \forall s \in \{1, \dots, \#S\}$ and an action \mathbf{a} , the goal of the P-step is to compute: $P(S_{I::\mathbf{a}}^s | O_t), \forall s \in \{1, \dots, \#S\}$

and $P(O_{I::\mathbf{a}}^o | O_t), \forall o \in \{1, \dots, \#O\}$. Note that by definition, we let $P(S_I^m | O_t) \triangleq P(S_t^m | O_t)$ if $I = \emptyset$. To derive the inference and prediction steps, the following sections make use of the sum-rule, product-rule, and d-separation criterion (c.f., Appendix C for details about those properties). Intuitively, d-separation refers to the use of conditional independencies (inherent in the Markov blankets of any random variable) to simplify the dependency structure and requisite computations of conditional probabilities.

2.3.1 INFERENCE STEP

As just stated, the goal of the I-step is to compute $P(S_t^m | O_t), \forall m \in \{1, \dots, \#S\}$. First, we re-write the posterior computation to fit the kind of problem that belief propagation — also known as the sum-product algorithm — can solve:

$$\begin{aligned}
 P(S_t^m | O_t) &\propto P(S_t^m, O_t) && \text{(Bayes theorem)} \\
 &= \sum_{\sim S_t^m} P(S_t, O_t) && \text{(sum rule)} \\
 &= \sum_{\sim S_t^m} \prod_{o=1}^{\#O} P(O_t^o | \rho_t^o) \prod_{s=1}^{\#S} P(S_t^s) && \text{(product rule \& d-separation)}
 \end{aligned}$$

where $S_t = \{S_t^s \mid s = 1, \dots, \#S\}$ is the set of all latent states at time t , $\sim S_t^m = S_t \setminus S_t^m$ is the set of all latent states at time t except S_t^m , and the summation is over all possible configurations of $\sim S_t^m$, i.e., we are marginalizing out all states, apart from one; thus $P(S_t, O_t)$ has $\#S + \#O$ dimensions, while $P(S_t^m, O_t)$ has $1 + \#O$ dimensions. Since $\rho_t^o \subseteq S_t$, the expression inside the summation is a function $g(S_t)$ that factorizes as follows:

$$\begin{aligned}
 g(S_t) &= \prod_{o=1}^{\#O} P(O_t^o | \rho_t^o) \prod_{s=1}^{\#S} P(S_t^s) \\
 &\triangleq \prod_{i=1}^N f_i(X_i),
 \end{aligned}$$

where $X_i \subseteq S_t$ for all $i \in \{1, \dots, \#O + \#S\}$, the number of factors is $N = \#O + \#S$, and:

$$f_i(X_i) \triangleq \begin{cases} P(O_t^i | \rho_t^i) & \text{if } i \in \{1, \dots, \#O\} \\ P(S_t^{i-\#O}) & \text{if } i \in \{\#O + 1, \dots, \#N\} \end{cases},$$

where $\#N = \#O + \#S$. Note that, because O_t^o (denoted O_t^i here) are known constants, we do not specify that $g(S_t)$ depends on O_t^o . To conclude, by substituting the definition of $g(S_t)$ into the formula of the posterior $P(S_t^m | O_t)$

presented above, we get:

$$P(S_t^m | O_t) \propto \sum_{\sim S_t^m} g(S_t),$$

which means that the posterior $P(S_t^m | O_t)$ can be computed by first marginalizing $g(S_t)$ w.r.t. S_t^m , i.e.,

$$g(S_t^m) = \sum_{\sim S_t^m} g(S_t),$$

and then normalizing:

$$P(S_t^m | O_t) = \frac{g(S_t^m)}{\sum_{S_t^m} g(S_t^m)}.$$

The marginalization of $g(S_t)$ can be performed efficiently using belief propagation (Kschischang et al., 2001), which can be understood as a message passing algorithm on a factor graph. The message from a node x to a factor f is given by:

$$m_{x \rightarrow f}(x) = \prod_{h \in n(x) \setminus \{f\}} m_{h \rightarrow x}(x),$$

where $n(x)$ are the neighbours of x in the factor graph. Note, in a factor graph the neighbours of a random variable are factors. Moreover, the message from a factor f to a node x is given by:

$$m_{f \rightarrow x}(x) = \sum_Y \left(f(X) \prod_{y \in Y} m_{y \rightarrow f}(y) \right),$$

where $X = n(f)$ are the neighbours of f in the factor graph, $Y = X \setminus \{x\}$ are all the neighbours of f except x , and the summation is over all possible configurations of the variables in Y . Note, in a factor graph the neighbours of a factor are random variables. Once all the messages have been computed, the marginalization of $g(S_t)$ w.r.t. S_t^m is given by the product of all the incoming messages of the node S_t^m , i.e.,

$$g(S_t^m) = \prod_{f \in n(S_t^m)} m_{f \rightarrow S_t^m}(S_t^m).$$

2.3.2 PREDICTION STEP

The P-step is analogous to the prediction step of Bayesian filtering (Fox et al., 2003). Given $P(S_I^s | O_t)$ for each $s \in \{1, \dots, \#S\}$ and an action \mathbf{a} , the goal of the P-step is to compute $P(S_{I::\mathbf{a}}^s | O_t)$ for each latent state $s \in \{1, \dots, \#S\}$ and $P(O_{I::\mathbf{a}}^o | O_t)$ for each future observation

$o \in \{1, \dots, \#O\}$. For the sake of brevity, we let $J \triangleq I::\mathbf{a}$. Let's start with the computation of $P(S_{I::\mathbf{a}}^s | O_t)$:

$$\begin{aligned} P(S_{I::\mathbf{a}}^s | O_t) &\triangleq P(S_J^s | O_t) \\ &= \sum_{\rho_J^s} P(S_J^s, \rho_J^s | O_t) && \text{(sum rule)} \\ &= \sum_{\rho_J^s} P(S_J^s | \rho_J^s, O_t) P(\rho_J^s | O_t) && \text{(product rule)} \\ &= \sum_{\rho_J^s} P(S_J^s | \rho_J^s) P(\rho_J^s | O_t) && \text{(d-separation)} \\ &\approx \sum_{\rho_J^s} P(S_J^s | \rho_J^s) \prod_{i=1}^{\#\rho_J^s} P(\rho_{J,i}^s | O_t) && \text{(mean-field approximation)} \end{aligned}$$

where $\#\rho_J^s$ is the number of parents of S_J^s , and $\rho_{J,i}^s$ is the i -th parent of S_J^s . Importantly, $P(S_J^s | \rho_J^s)$ is known from the definition of the generative model. Moreover, since $\rho_{J,i}^s \in S_I$, then $P(\rho_{J,i}^s | O_t) = P(S_I^m | O_t)$ for some $m \in \{1, \dots, \#S\}$. Thus, $P(\rho_{J,i}^s | O_t)$ is given as input to the P-step, i.e., $P(\rho_{J,i}^s | O_t)$ is a known distribution. Similarly, the computation of $P(O_{I::\mathbf{a}}^o | O_t)$ proceeds as follows:

$$\begin{aligned} P(O_{I::\mathbf{a}}^o | O_t) &\triangleq P(O_J^o | O_t) \\ &= \sum_{\rho_J^o} P(O_J^o, \rho_J^o | O_t) && \text{(sum rule)} \\ &= \sum_{\rho_J^o} P(O_J^o | \rho_J^o, O_t) P(\rho_J^o | O_t) && \text{(product rule)} \\ &= \sum_{\rho_J^o} P(O_J^o | \rho_J^o) P(\rho_J^o | O_t) && \text{(d-separation)} \\ &\approx \sum_{\rho_J^o} P(O_J^o | \rho_J^o) \prod_{i=1}^{\#\rho_J^o} P(\rho_{J,i}^o | O_t) && \text{(mean-field approximation)} \end{aligned}$$

where $\#\rho_J^o$ is the number of parents of O_J^o , and $\rho_{J,i}^o$ is the i -th parent of O_J^o . Importantly, $P(O_J^o | \rho_J^o)$ is known from the definition of the generative model. Moreover, since $\rho_{J,i}^o \in S_J$, then $P(\rho_{J,i}^o | O_t) = P(S_J^s | O_t)$ for some $s \in \{1, \dots, \#S\}$. Thus, $P(\rho_{J,i}^o | O_t)$ has already been computed during the first stage of the P-step and is a known distribution, c.f., derivation of $P(S_{I::\mathbf{a}}^s | O_t) \triangleq P(S_J^s | O_t)$.

2.4 Expected Free Energy

In this section, we discuss the definition of the expected free energy, which quantifies the cost of pursuing a par-

ticular sequence of actions that underwrite planning, cf. Section 2.5. The expected free energy (see below) is composed of the risk and ambiguity terms. The risk terms quantify how much the posterior beliefs over future observations (computed by the P-step) diverge from the prior preferences of the agent. On the other hand, the ambiguity terms correspond to the expected uncertainty of the likelihood mapping, where the expectation is with respect to the posterior beliefs over states computed by the P-step.

First, we partition the set of observations $O_I = \{O_I^o \mid o = 1, \dots, \#O\}$ into disjoint subsets X_i^I , i.e., $O_I = X_1^I \cup \dots \cup X_N^I$ and $X_i^I \cap X_j^I = \emptyset$ if $i \neq j$. Then, we define the prior preferences over the i -th subset of observations as: $V(X_i^I) = \text{Cat}(\mathbf{C}^i)$. This formulation allows us to define prior preferences over subsets of random variables, and will be useful in Section 3.1, where the agent needs to possess preferences that depend upon both the shape and (X, Y) position of the object. Finally, the expected free energy, which needs to be minimised, is given by:

$$\mathbf{G}_I \triangleq \sum_{i=1}^N \left(\underbrace{D_{\text{KL}}[P(X_i^I|O_t) \parallel V(X_i^I)]}_{\text{risk of } i\text{-th set of observations}} \right) + \sum_{o=1}^{\#O} \left(\underbrace{\mathbb{E}_{P(\rho_I^o|O_t)}[\mathbf{H}[P(O_I^o|\rho_I^o)]]}_{\text{ambiguity of } o\text{-th observation}} \right), \quad (1)$$

where $P(X_i^I|O_t)$ and $P(\rho_I^o|O_t)$ are the posteriors over the i -th subset of observations and the parent of O_I^o , respectively, and $P(O_I^o|\rho_I^o)$ is known from the generative model. Assuming a mean-field approximation, those posteriors are given by:

$$P(\rho_I^o|O_t) \approx \prod_{i=1}^{\#\rho_I^o} P(\rho_{I,i}^o|O_t)$$

$$P(X_i^I|O_t) \approx \prod_{O_I^o \in X_i} P(O_I^o|O_t)$$

where $P(O_I^o|O_t)$ and $P(\rho_{I,i}^o|O_t)$ are the posteriors over O_I^o and the i -th parent of O_I^o , respectively. Note, both $P(O_I^o|O_t)$ and $P(\rho_{I,i}^o|O_t)$ were computed during the P-step. The definition of the expected free energy given by (1) may not be very intuitive. Fortunately, the special case where each subset contains a single observation, i.e.,

$X_o^I = O_I^o$, leads to the following equation:

$$\mathbf{G}_I \triangleq \sum_{o=1}^{\#O} \left(\underbrace{D_{\text{KL}}[P(O_I^o|O_t) \parallel V(O_I^o)]}_{\text{risk of } o\text{-th observation}} + \underbrace{\mathbb{E}_{P(\rho_I^o|O_t)}[\mathbf{H}[P(O_I^o|\rho_I^o)]]}_{\text{ambiguity of } o\text{-th observation}} \right),$$

which is the summation over all observations O_I^o of the expected free energy of O_I^o , i.e., the risk of O_I^o plus the ambiguity of O_I^o . Finally, our framework allows to specify prior preferences over only a subset of variables in O_I . For example, if a task contains four variables, i.e., $O_I^x, O_I^y, O_I^{shape}$ and O_I^{scale} , but it only makes sense to have preferences over three of them, i.e., O_I^x, O_I^y and O_I^{shape} , then the prior preference over the fourth variable is set to the posterior over this random variable, i.e., $V(O_I^{scale}) \triangleq P(O_I^{scale}|O_t)$. In other words, not having prior preferences over a random variable is viewed by our framework as liking whatever we predict will happen. Effectively, this renders the risk term associated with such variable equal to zero, i.e.,

$$D_{\text{KL}}[P(O_I^{scale}|O_t) \parallel V(O_I^{scale})] = D_{\text{KL}}[P(O_I^{scale}|O_t) \parallel P(O_I^{scale}|O_t)] = 0.$$

2.5 Planning: the MCTS algorithm

In this section, we describe the planning algorithm used by *BTAI_{3MF}*. At the beginning of a trial when $t = 0$, the agent is provided with the initial observations O_0 . The I-step is performed and returns the posterior over all latent states, i.e., $P(S_0^s|O_0)$ for all $s \in \{1, \dots, \#S\}$, according to the prior over the initial hidden states provided by the modeller, i.e., $P(S_0^s)$ for all $s \in \{1, \dots, \#S\}$, and the available observations O_0 .

Then, we use the UCT criterion to determine which node in the tree should be expanded. Let the tree's root TS_t be called the current node. If the current node has no children, then it is selected for expansion. Alternatively, the child with the highest UCT criterion becomes the new current node and the process is iterated until we reach a leaf node (i.e. a node from which no action has previously been selected). The UCT criterion (Browne et al., 2012) for the j -th child of the current node is given by:

$$UCT_j = -\bar{\mathbf{G}}_j + C_{\text{explore}} \sqrt{\frac{\ln n}{n_j}}, \quad (2)$$

where $\bar{\mathbf{G}}_j$ is the average expected free energy calculated with respect to the actions selected from the j -th child, C_{explore} is the exploration constant that modulates the

amount of exploration at the tree level, n is the number of times the current node has been visited, and n_j is the number of times the j -th child has been visited.

Let S_I be the (leaf) node selected by the above selection procedure. We then expand all the children of S_I , i.e., all the states of the form $S_{I::\mathbf{a}}$, where $\mathbf{a} \in \{1, \dots, \#A\}$ is an arbitrary action, $\#A$ is the number of available actions, and $I::\mathbf{a}$ is the multi-index obtained by appending the action \mathbf{a} at the end of the sequence defined by I . Next, we perform the P-step for each action \mathbf{a} , and obtain $P(S_{I::\mathbf{a}}^s | O_t)$ for each latent state $s \in \{1, \dots, \#S\}$ and $P(O_{I::\mathbf{a}}^o | O_t)$ for each future observation $o \in \{1, \dots, \#O\}$.

Then, we need to estimate the cost of (virtually) taking each possible action. The cost in active inference is the expected free energy given by (1). Next, we assume that the agent will always perform the action with the lowest cost, and back-propagate the cost of the best (virtual) action toward the root of the tree. Formally, we write the update as follows:

$$\forall K \in \mathbb{A}_I \cup \{I\}, \quad \mathbf{G}_K \leftarrow \mathbf{G}_K + \min_{\mathbf{a} \in \{1, \dots, \#A\}} \mathbf{G}_{I::\mathbf{a}}, \quad (3)$$

where I is the multi-index of the node that was selected for (virtual) expansion, and \mathbb{A}_I is the set of all multi-indices corresponding to ancestors of TS_I . During the back propagation, we also update the number of visits as follows:

$$\forall K \in \mathbb{A}_I \cup \{I\}, \quad n_K \leftarrow n_K + 1. \quad (4)$$

If we let \mathbf{G}_K^{aggr} be the aggregated cost of an arbitrary node S_K obtained by applying Equation 3 after each expansion, then we are now able to express $\bar{\mathbf{G}}_K$ formally as:

$$\bar{\mathbf{G}}_K = \frac{\mathbf{G}_K^{aggr}}{n_K}.$$

The planning procedure described above ends when the maximum number of planning iterations is reached.

2.6 Action selection

After performing planning, the agent needs to choose the action to perform in the environment. As discussed in Section 3.1 of (Browne et al., 2012), many possible mechanisms can be used to select the action to perform in the environment. $BTAI_{3MF}$ performs the action corresponding to the root child with the highest number of visits. Formally, this is expressed as:

$$\mathbf{a}^* = \arg \max_{\mathbf{a} \in \{1, \dots, \#A\}} n_{(\mathbf{a})}, \quad (5)$$

where \mathbf{a}^* is the action performed in the environment, and $n_{(\mathbf{a})}$ is the number of visits of the root child corresponding to action \mathbf{a} .

2.7 Closing the action-perception cycle

After performing an action \mathbf{a}^* in the environment, the agent receives a new observation O_{t+1} , and needs to use this observation to compute the posterior over the latent states at time $t + 1$, i.e., $P(S_{t+1}^s | O_{t+1})$ for all $s \in \{1, \dots, \#S\}$. This can be achieved by performing the I-step, but requires the agent to have prior beliefs over the latent states at time $t + 1$, i.e., $P(S_{t+1}^s)$ for all $s \in \{1, \dots, \#S\}$, in addition to the new observation O_{t+1} obtained from the environment. In this paper, we define those prior beliefs as:

$$P(S_{t+1}^s) = P(S_I^s | O_t), \text{ for all } s \in \{1, \dots, \#S\},$$

where $I = (\mathbf{a}^*)$ is a sequence of actions containing the action \mathbf{a}^* performed in the environment, $P(S_I^s | O_t)$ is the predictive posterior computed by the P-step when assuming that action \mathbf{a}^* is performed. In other words, the predictive posterior $P(S_I^s | O_t)$ computed by the P-step at time t , is used as an empirical prior $P(S_{t+1}^s)$ at time $t + 1$. This empirical prior $P(S_{t+1}^s)$ along with the new observation O_{t+1} can then be used to compute the posterior $P(S_{t+1}^s | O_{t+1})$ for all $s \in \{1, \dots, \#S\}$. This posterior will be used to perform planning in the next action-perception cycle. Appendix A provides an algorithm summarizing our approach, c.f., Algorithm 1.

3. Results

In this section, we compare our new approach to BTAI with variational message passing ($BTAI_{VMP}$) and BTAI with Bayesian filtering ($BTAI_{BF}$). Section 3.1 presents the simplified version of the dSprites environment on which the agents are compared. Section 3.2 describes how the task is modelled by the $BTAI_{VMP}$ agent and reports its performance, finally, Sections 3.3 and 3.4 do the same for the $BTAI_{BF}$ and $BTAI_{3MF}$ agents. For the reader interested in implementing a custom $BTAI_{3MF}$ agent, Appendix A provides a tutorial of how to create such an agent using our framework, and Appendix B describes a graphical user interface (GUI) that can be used to inspect the model. This GUI displays the structure of the generative model and prior preferences, the posterior beliefs of each latent variable, the messages sent throughout the factor graph to perform inference, the information related to the MCTS algorithm, and the expected free energy (EFE) of each node in the future. It also shows how the EFE decomposes into the risk and ambiguity terms.

3.1 dSprites Environment

The dSprites environment (Fountas et al., 2020) is based on the dSprites dataset (Matthey et al., 2017) initially designed for analysing the latent representation learned by variational auto-encoders (Doersch, 2016). The dSprites dataset is composed of images of squares, ellipses and hearts. Each image contains one shape (square, ellipse or heart) with its own scale, orientation, and (X, Y) position. In the dSprites environment, the agent is able to move those shapes around by performing four actions (i.e., UP, DOWN, LEFT, RIGHT). To make planning tractable, the action selected by the agent is executed eight times in the environment before the beginning of the next action-perception cycle, i.e., the X or Y position is increased or decreased by eight between time step t and $t+1$. The goal of the agent is to move all squares towards the bottom-left corner of the image and all ellipses and hearts towards the bottom-right corner of the image, c.f. Figure 4.



Figure 4: This figure illustrates the dSprites environment, in which the agent must move all squares towards the bottom-left corner of the image and all ellipses and hearts towards the bottom-right corner of the image. The red arrows show the behaviour expected from the agent.




Figure 5: This figure illustrates the observations made by the agent when using a coarse-grained representation with a granularity of eight on the input image. On the left, one can see an image from the dSprites dataset and a grid containing red squares of 8×8 pixels. Any positions in those 8×8 squares are indistinguishable from the perspective of the agent. Also, the bottom most row is an absorbing row used to specify the prior preferences of the agent, i.e. the green square is the goal state and the orange squares correspond to undesirable states. Finally, the three tables on the right contain the indices observed by the $BTAI_{VMP}$ and $BTAI_{BF}$ agents for each type of shape at each possible position.

Since BTAI is a tabular model whose likelihood and transition mappings are represented using matrices, the agent does not directly take images as inputs. Instead, the metadata of the dSprites dataset is used to specify the state space. In particular, the agent observes the type of shape (i.e., square, ellipse, or heart), the scale and orientation of the shape, as well as a coarse-grained version of the shape’s true position. Importantly, the original images are composed of 32 possible values for both the X and Y positions of the shapes. A coarse-grained representation with a granularity of two means that the agent is only able to perceive 16×16 images, and thus, the positions at coordinate $(0, 0)$, $(0, 1)$, $(1, 0)$ and $(1, 1)$ are indistinguishable. Figure 5 illustrates the coarse grained representation with a granularity of eight and the corresponding indices observed by the $BTAI_{VMP}$ and $BTAI_{BF}$ agents. Note that this modification of the observation space can be seen as a form of state aggregation (Ren and Krogh, 2002). Finally, as shown in Figure 5, the prior preferences of the agent are specified over an absorbing row below the dSprites image. This absorbing row ensures that the agent selects the action “down” when standing in the “appropriate corner”, i.e., bottom-left corner for squares and bottom-right corner for ellipses and hearts.

The evaluation of the agent’s performance is based on the reward obtained by the agent. Briefly, the agent receives a reward of -1 , if it never enters the absorbing row or if it does so at the antipode of the appropriate corner. As the agent enters the absorbing row closer and closer to the appropriate corner, its reward increases until reaching

a maximum of 1. The percentage of the task solved (i.e., the evaluation metric) is calculated as follows:

$$P(\text{solved}) = \frac{\text{total rewards} + \text{number of runs}}{2.0 \times \text{number of runs}}.$$

Intuitively, the numerator shifts the rewards so that they are bounded between zero and two, and the denominator renormalises the reward to give a score between zero and one. A score of zero therefore corresponds to an agent always failing to enter the absorbing row or doing so at the antipode of the appropriate corner. In contrast, a score of one corresponds to an agent always entering the absorbing row through the appropriate corner.

3.2 $BTAI_{VMP}$ modeling approach and results

In this section, we evaluate $BTAI_{VMP}$ (Champion et al., 2022b,a) on the dSprites environment. As shown in Figure 5, $BTAI_{VMP}$ observes one index for each possible configuration of shape, and (X, Y) positions. Importantly, this version of BTAI suffers from the exponential growth described in the introduction, and thus does not model the scale and orientation modalities. Also, to make the inference and planning process tractable, the granularity of the coarse-grained representation was set to four or eight. Table 1 provides the value of each hyper-parameter used by $BTAI_{VMP}$ in this section. Note, the hyper-parameter values are the same for all BTAI models presented in this paper. Only the number of action perception cycles, and the number of planning iterations may vary from one experiment to the next.

Briefly, the agent is able to solve 88.5% of the task when using a granularity of eight, c.f. Table 2. To understand why $BTAI_{VMP}$ was not able to solve the task with 100% accuracy, let us consider the example of an ellipse at position $(24, 31)$. With a granularity of eight, the agent perceives that the ellipse is in the bottom-right corner of the image, i.e., in the red square just above the goal state in Figure 5. From the agent’s perspective, it is thus optimal to pick the action “down” to reach the goal state. However, in reality, the agent will not receive the maximum reward because its true X position is 24 instead of the optimal X position of 31.

Name	Value
NB_SIMULATIONS	100
NB_ACTION_PERCEPTION_CYCLES	30
NB_PLANNING_STEPS	10, 25 or 50
EXPLORATION_CONSTANT	2.4
PRECISION_PRIOR_PREFERENCES	2
PRECISION_ACTION_SELECTION	100
EVALUATION_TYPE	EFE

Table 1: The value of each hyper-parameter used by $BTAI_{VMP}$ in this section. NB_SIMULATIONS is the number of simulations run during the experiment. NB_ACTION_PERCEPTION_CYCLES is the maximum number of actions executed in each simulation, after which the simulation is terminated. NB_PLANNING_STEPS is the number of planning iterations performed by the agent. EXPLORATION_CONSTANT is the exploration constant of the UCT criterion. PRECISION_PRIOR_PREFERENCES is the precision of the prior preferences. PRECISION_ACTION_SELECTION is the precision of the distribution used for action selection. EVALUATION_TYPE is the type of cost used to evaluate the node during the tree search. Those hyper-parameters can be used to re-run the experiments using the code of the following GitHub repository: https://github.com/ChampiB/Experiments_AI_TS.

Planning iterations	P(solved)	Time (sec)
10	0.813	0.859 ± 0.868
25	0.846	0.862 ± 0.958
50	0.885	1.286 ± 1.261

Table 2: The percentage of the dSprites environment solved by the $BTAI_{VMP}$ agent when using a granularity of eight, c.f. Figure 5. The last column reports the average execution time required for one simulation and the associated standard deviation.

As shown in Table 3, we can improve the agent’s performance, by using a granularity of four. This allows the agent to differentiate between a larger number of (X, Y) positions, i.e., it reduces the size of the red square in Figure 5. With this setting, the agent is able to solve 96.9% of the task. However, when decreasing the granularity, the number of states goes up, and so does the width and height of the \mathbf{A} and \mathbf{B} matrices. As a result, more memory and computational time is required for the inference and planning process. This highlights a trade-off between the agent’s performance and the amount of memory and time required. Indeed, a smaller granularity leads to better performance, but requires more time and memory.

Planning iterations	P(solved)	Time (sec)
10	0.859	3.957 ± 4.027
25	0.933	3.711 ± 4.625
50	0.969	5.107 ± 5.337

Table 3: The percentage of the dSprites environment solved by the $BTAI_{VMP}$ agent when using a granularity of four. In this setting, there are $9 \times 8 \times 3 = 216$ states. The last column reports the average execution time required for one simulation and the associated standard deviation.

3.3 $BTAI_{BF}$ modeling approach and results

In this section, we evaluate $BTAI_{BF}$ (Champion et al., 2021a) on the dSprites environment. The Bayesian filtering version is similar to the variational message passing but uses belief propagation to furnish updates that correspond to exact Bayesian inference. In the next results section, we look at the mean field version, in which the same kind of belief propagation scheme is used but under a mean field approximation that greatly simplifies the computation of expected free energy and the requisite tree searches over plausible policies. As shown in Figure 5, $BTAI_{BF}$ observes one index for each possible configuration of shape, and (X, Y) positions. Also, to make the inference and planning process tractable, the granularity of the coarse-grained representation was set to two, four or eight. Table 4 provides the value of each hyper-parameter used by $BTAI_{BF}$ in this section. Note, the hyper-parameter values are the same for all BTAI models presented in this paper. Only the number of action perception cycles, and the number of planning iterations may vary from one experiment to the next.

As shown in Table 5, the agent is able to solve: 86.1% of the task when using a granularity of eight, 97.7% of the task when using a granularity of four, and 98.6% of the task when using a granularity of two. However, as the performance improves from 86.1% to 98.6%, the computational time required to run each simulation skyrockets from around 50 milliseconds to around 17.5 seconds. In other words, a simulation with a granularity of two is 350 times slower than a simulation with a granularity of eight.

Name	Value
NB_SIMULATIONS	100
NB_ACTION_PERCEPTION_CYCLES	20
NB_PLANNING_STEPS	50
EXPLORATION_CONSTANT	2.4
PRECISION_PRIOR_PREFERENCES	1
PRECISION_ACTION_SELECTION	100
EVALUATION_TYPE	EFE

Table 4: The value of each hyper-parameter used by $BTAI_{3MF}$ in this section. Please refer to Table 1 for a detailed description of the hyper-parameters. Those hyper-parameters can be used to re-run the experiments using the code of the following GitHub repository: https://github.com/ChampiB/BTAI_3MF.

PI	Granularity	P(solved)	Time (ms)
50	8	0.861	49.93 ± 36.4124
50	4	0.977	241.63 ± 118.379
50	2	0.986	17503.8 ± 12882.8

Table 5: The percentage of the dSprites environment solved by the $BTAI_{BF}$ agent when using a granularity of eight, four and two. Note, when a granularity of two is used, there are $17 \times 16 \times 3 = 816$ possible states. PI stands for planning iterations. The last column reports the average execution time required for one simulation and the associated standard deviation. Note, the change in time granularity to milliseconds.

3.4 $BTAI_{3MF}$ modeling approach and results

In this section, we evaluate our new approach ($BTAI_{3MF}$) on the dSprites environment. The mean field approximation is similar to the Bayesian filtering scheme; however, it is much more efficient (and expressive) because it, effectively, allows one to replace very high dimensional tensors (encoding the dependencies of outcomes on multiple latent states) with a set of matrices that consider each latent state separately. In contrast to what is shown in Figure 5, $BTAI_{3MF}$ does not observe one index for each possible configuration of shape, and (X, Y) positions. Instead, $BTAI_{3MF}$ has five observed variables representing the shape, the orientation, the scale, as well as the X and Y position, respectively. Each of those observed variable has its hidden state counterparts. Each observation depends on its hidden state counterparts through an identity matrix. This parametrisation is common in the literature on active inference, see (Sajid et al., 2021) for an example. The transition mappings of the hidden variables rep-

representing the shape, orientation, and scale, are defined as an identity matrix. This forwards the state value at time t to the next time step $t + 1$. For the hidden variables representing the X and Y position of the shape, the transition is set to reflect the dynamics of the dSprites environment when the actions taken are repeated eight times, i.e., if the action “DOWN” is selected, then the agent’s position in Y will be decreased by eight before the start of the next action-perception cycle (Fountas et al., 2020).

The hyper-parameters used in those simulations are presented in Table 6. Note, the hyper-parameter values are the same for all BTAI models presented in this paper. Only the number of action perception cycles, and the number of planning iterations may vary from one experiment to the next.

Table 7 shows the results obtained by $BTAI_{3MF}$ on the dSprites environment when running 100 trials. Due to the change in the format of representations, the agent exhibits little increase in execution time as the granularity decreases, however, in general, the capacity to solve the task increases with this reduction in granularity. When a granularity of one is used, the agent is able to solve the task perfectly with 150 planning iterations.

Note, the agent using a granularity of 1 and 150 planning iterations is as fast as the agent using a granularity of 1 and 50 planning iterations. This is because as the number of planning iterations increases the agent requires more computation time per action-perception cycle, but as the agent performance increases on the task, the agent reaches the goal state faster, and therefore requires less action-perception cycles per simulation. To conclude, the agent with 150 planning iterations requires less action-perception cycles per simulation, but more time per action-perception cycle than the agent with 50 planning iterations. The code relevant to this section is available at the following URL: https://github.com/ChampiB/BTAI_3MF.

Name	Value
NB_SIMULATIONS	100
NB_ACTION_PERCEPTION_CYCLES	50
NB_PLANNING_STEPS	50/100/150
EXPLORATION_CONSTANT	2.4
PRECISION_PRIOR_PREFERENCES	1
EVALUATION_TYPE	EFE

Table 6: The value of each hyper-parameter used by $BTAI_{3MF}$ in this section. Please refer to Table 1 for a detailed description of the hyper-parameters. Those hyper-parameters can be used to re-run the experiments using the code of the following GitHub repository: https://github.com/ChampiB/BTAI_3MF.

PI	Granularity	P(solved)	Time (sec)
50	8	0.895	1.279 ± 12.8
50	4	0.977	1.279 ± 12.8
50	2	0.996	1.279 ± 12.8
50	1	0.72	2.559 ± 18.01
100	1	0.77	5.119 ± 25.209
150	1	1	2.559 ± 18.01

Table 7: This table presents the percentage of the dSprites environment solved by the $BTAI_{3MF}$ agent when using a granularity of eight, four, two and one. Note, when a granularity of one is used, there are $33 \times 32 \times 3 \times 40 \times 6 = 760,320$ possible state configurations. PI stands for planning iterations. The last column reports the average execution time required of one simulation and the associated standard deviation.

4. Conclusion

In this paper, we presented a new version of Branching Time Active Inference that allows for modelling of several observed and latent variables. Taken together, those variables constitute a temporal slice. Within a slice, the model is equipped with prior beliefs over the initial latent variables, and each observation depends on a subset of the latent variables through the likelihood mapping. Additionally, the latent states evolve over time according to the transition mapping that describes how each latent variable at time $t + 1$ is generated from a subset of the hidden states at time t and the action taken.

At the beginning of each trial, the agent makes an observation for each observed variable, and computes the posterior over the latent variables using belief propagation. Then, a Monte-Carlo tree search is performed to explore the space of possible policies. During the tree

search, each planning iteration starts by selecting a node to expand using the UCT criterion. Then, the children of the selected node are expanded, i.e., one child per action. Next, the posterior over the latent variables of the expanded nodes is computed by performing forward predictions using the known transition mapping, and the posterior beliefs over the latent states of the node selected for expansion. Once the posterior is computed, the expected free energy can be computed and back-propagated through the tree. The planning process stops after reaching a maximum number of iterations.

In the results section, we compared our new approach, called $BT AI_{3MF}$, to two earlier versions of branching time active inference, named $BT AI_{VMP}$ (Champion et al., 2022b,a) and $BT AI_{BF}$ (Champion et al., 2021a). Briefly, at the current time step t : $BT AI_{VMP}$ performs variational message passing (VMP) with a variational distribution composed of only one factor, $BT AI_{BF}$ performs exact inference using Bayes theorem, and $BT AI_{3MF}$ implements belief propagation to compute the marginal posterior over each latent variable. For the hidden variables in the future, $BT AI_{VMP}$ does the same mean-field approximation as at time step t and performs VMP, $BT AI_{BF}$ performs Bayesian prediction to compute the posterior over the only latent variable being modelled, and likewise, $BT AI_{3MF}$ performs prediction to compute the posterior over all future latent variables.

Since, none of the aforementioned approaches are equipped with deep neural networks, we compared them on a version of the dSprites environment in which the metadata of the dSprites dataset are used as inputs to the model instead of the dSprites images. The best performance obtained by $BT AI_{VMP}$ was to solve 96.9% of the task in 5.1 seconds. Importantly, $BT AI_{VMP}$ was previously compared to active inference as implemented in SPM both theoretically and experimentally (Champion et al., 2022b,a). $BT AI_{BF}$ was able to solve 98.6% of the task but at the cost of 17.5 seconds of computation. Note, $BT AI_{BF}$ was using a granularity of two (i.e., 816 states) while $BT AI_{VMP}$ was using a granularity of four (i.e., 216 states), which is why $BT AI_{BF}$ seems to be three times slower than $BT AI_{VMP}$. In reality, if $BT AI_{BF}$ had been using a granularity of four, it would have been much faster than $BT AI_{VMP}$ while maintaining a similar performance, i.e., around 96.9% of the task solved. Finally, $BT AI_{3MF}$ outperformed both of its predecessors by solving the task completely (100%, granularity of 1) in only 2.559 seconds. Importantly, $BT AI_{3MF}$ was able to model all the modalities of the dSprites environment for a total of 760,320 possible states.

In addition to the major boost in performance and computational time, $BT AI_{3MF}$ provides an improved modelling capacity. Indeed, the framework can now handle the modelling of several observed and latent variables, and takes advantage of the factorisation of the generative model to perform inference efficiently. As described in detail in Appendix A, we also provide a high level notation for the creation of $BT AI_{3MF}$ that aims to make our approach as straightforward as possible to apply to new domains. The high-level notational language allows the user to create models by simply declaring the variables it contains, and the dependencies among those variables. Then, the framework performs the inference process automatically. Moreover, driven by the need for interpretability, we developed a graphical user interface to analyse the behaviour and reasoning of our agent, which is described in Appendix B.

There are two major directions of future research that may be explored to keep scaling up this framework. First, $BT AI_{3MF}$ is not yet equipped with deep neural networks (DNNs), and is therefore unable to handle certain types of inputs, such as images. In addition to the integration of DNNs into the framework, further research should be performed in order to learn useful sequences of actions. Typically, in the current version of $BT AI_{3MF}$, we built in the fact that each action should be repeated eight times in a row. This inductive bias works well in the context of the dSprites environment, but may be a limitation in other contexts.

It is also worth reflecting on how the $BT AI_{3MF}$ model sits with theories of brain function. In this respect, it is interesting to consider neural correlates of the “standard” approach that $BT AI_{3MF}$ is being placed in opposition to. As previously discussed, this standard active inference approach could be considered as monolithically tabular; that is, the key matrices, such as the likelihood mapping (the \mathbf{A} matrix) and the transition mapping (the \mathbf{B} matrix), grow in size exponentially with the number of states and observations. This is simply due to a combinatorial explosion, e.g. the set of all combinations of states grows intractably with the number of states.

How would the combinations of states in the monolithic tabular approach be represented in the brain? The obvious neural correlate would be conjunctive (binding) neurons (O’Reilly and Rudy, 2001), which become active when multiple feature values are present; for example, one might have a neural unit for every X, Y combination in the dSprites environment. If this is to be realised with a fully localist code, i.e. one unit for every combination, in the absence of any hierarchical structure, the required number of conjunctive units would explode in the same

way as the \mathbf{A} and \mathbf{B} matrices do. This is why some models have proposed a binding resource that supports distributed (rather than localist) representations (Bowman and Wyble, 2007), which scale more tractably.

$BTAI_{3MF}$ avoids this combinatorial explosion by not combining features, enabling them to be represented separately. In a very basic sense, this separated representation is consistent with the observation that the brain contains distinct, physically separated, feature maps, e.g. Itti et al. (1998). Thus, at least to some extent, different feature dimensions are processed separately in the brain, as they are in $BTAI_{3MF}$.

The time-slice idea in $BTAI_{3MF}$ assumes a kind of discrete synchronising global clock. Thus, even though features have been separated from one another and may be considered to execute in different parts of the system, they update in lock-step. That is, implicitly, time is a binder, it determines which values of different feature dimensions/states are associated, e.g. an X-dimension value is associated with a particular Y-dimension value because they are so assigned in the same temporal slice. In this sense, in $BTAI_{3MF}$, time synchronisation resolves the binding problem.

This aspect of $BTAI_{3MF}$ resonates with theories of binding based upon oscillatory synchrony (Uhlhaas et al., 2009). These theories suggest that different feature dimensions are bound by the corresponding neurons firing in synchrony relative to an ongoing oscillation, with that ongoing oscillation potentially playing the role of a global clock. Such oscillatory synchrony can be seen as a way to resolve the binding problem that does not require conjunctive units.

Conjunction error experiments, e.g. Botella et al. (2001), are also relevant here. In these experiments, participants make errors in associating multiple feature dimensions, perceiving illusory percepts, e.g. if a red K is presented before a blue A in a rapid serial visual presentation stream, in some cases, a red A and a blue K is perceived. These experiments firstly, re-emphasize that different feature dimensions are processed separately, as per $BTAI_{3MF}$: if feature dimensions were not separated, then conjunction errors could not happen. Additionally though, these experiments suggest that there is not a “perfect” synchronising global clock, since if there were, there would not be any conjunction errors even despite separation of feature dimensions. Generating such conjunction error patterns is an interesting topic for future $BTAI_{3MF}$ modelling work.

Acknowledgments

We would like to thank the reviewers for their valuable feedback, which greatly improved the quality of the present paper.

References

- J Botella, M Suero, and MI Barriopedro. A model of the formation of illusory conjunctions in the time domain. *Journal of experimental psychology. Human perception and performance*, 27(6):1452–1467, December 2001. ISSN 0096-1523. doi: 10.1037//0096-1523.27.6.1452. URL <https://doi.org/10.1037//0096-1523.27.6.1452>.
- Matthew Botvinick and Marc Toussaint. Planning as inference. *Trends in Cognitive Sciences*, 16(10):485 – 488, 2012. ISSN 1364-6613. doi: <https://doi.org/10.1016/j.tics.2012.08.006>.
- Howard Bowman and Brad Wyble. The simultaneous type, serial token model of temporal attention and working memory. *Psychological review*, 114(1):38, 2007.
- C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.
- Théophile Champion, Marek Grześ, and Howard Bowman. Branching Time Active Inference with Bayesian Filtering, 2021a.
- Théophile Champion, Marek Grześ, and Howard Bowman. Realizing Active Inference in Variational Message Passing: The Outcome-Blind Certainty Seeker. *Neural Computation*, 33(10):2762–2826, 09 2021b. ISSN 0899-7667. doi: 10.1162/neco_a.01422. URL https://doi.org/10.1162/neco_a_01422.
- Théophile Champion, Howard Bowman, and Marek Grześ. Branching time active inference: Empirical study and complexity class analysis. *Neural Networks*, 2022a. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2022.05.010>. URL <https://www.sciencedirect.com/science/article/pii/S0893608022001824>.
- Théophile Champion, Lancelot Da Costa, Howard Bowman, and Marek Grześ. Branching time active inference: The theory and its generality. *Neural Networks*, 151:295–316, 2022b. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2022>.

- 03.036. URL <https://www.sciencedirect.com/science/article/pii/S08933608022001149>.
- Lancelot Da Costa, Thomas Parr, Noor Sajid, Sebastian Veselic, Victorita Neacsu, and Karl Friston. Active inference on discrete state-spaces: a synthesis, 2020.
- Maell Cullen, Ben Davey, Karl J. Friston, and Rosalyn J. Moran. Active inference in openai gym: A paradigm for computational investigations into psychiatric illness. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 3(9):809 – 818, 2018. ISSN 2451-9022. doi: <https://doi.org/10.1016/j.bpsc.2018.06.010>. URL <http://www.sciencedirect.com/science/article/pii/S2451902218301617>. Computational Methods and Modeling in Psychiatry.
- Carl Doersch. Tutorial on variational autoencoders, 2016.
- Thomas H. B. FitzGerald, Raymond J. Dolan, and Karl Friston. Dopamine, reward learning, and active inference. *Frontiers in Computational Neuroscience*, 9: 136, 2015. ISSN 1662-5188. doi: 10.3389/fncom.2015.00136. URL <https://www.frontiersin.org/article/10.3389/fncom.2015.00136>.
- Zafeirios Fountas, Noor Sajid, Pedro A. M. Mediano, and Karl Friston. Deep active inference agents using Monte-Carlo methods, 2020.
- V. Fox, J. Hightower, Lin Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *IEEE Pervasive Computing*, 2(3):24–33, 2003. doi: 10.1109/MPRV.2003.1228524.
- Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, John O Doherty, and Giovanni Pezzulo. Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68:862 – 879, 2016. ISSN 0149-7634. doi: <https://doi.org/10.1016/j.neubiorev.2016.06.022>.
- Karl J. Friston, Thomas Parr, and Bert de Vries. The graphical brain: Belief propagation and active inference. *Network Neuroscience*, 1(4):381–414, 2017. doi: 10.1162/NETN_a_00018. URL https://doi.org/10.1162/NETN_a_00018.
- L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259, 1998. doi: 10.1109/34.730558.
- Laurent Itti and Pierre Baldi. Bayesian surprise attracts human attention. *Vision Research*, 49(10):1295 – 1306, 2009. ISSN 0042-6989. doi: <https://doi.org/10.1016/j.visres.2008.09.007>. URL <http://www.sciencedirect.com/science/article/pii/S0042698908004380>. Visual Attention: Psychophysics, electrophysiology and neuroimaging.
- Frank R Kschischang, Brendan J Frey, and H-A Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on information theory*, 47(2):498–519, 2001.
- Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- Beren Millidge. Combining active inference and hierarchical predictive coding: A tutorial introduction and case study., 2019. URL <https://doi.org/10.31234/osf.io/kf6wc>.
- Randall C O’Reilly and Jerry W Rudy. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychological review*, 108(2):311, 2001.
- Corrado Pezzato, Carlos Hernandez, and Martijn Wisse. Active inference and behavior trees for reactive action planning and execution in robotics, 2020.
- Zhiyuan Ren and B.H. Krogh. State aggregation in Markov decision processes. In *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, volume 4, pages 3819–3824 vol.4, 2002. doi: 10.1109/CDC.2002.1184960.
- Noor Sajid, Philip J. Ball, Thomas Parr, and Karl J. Friston. Active Inference: Demystified and Compared. *Neural Computation*, 33(3):674–712, 03 2021. ISSN 0899-7667. doi: 10.1162/neco_a.01357. URL https://doi.org/10.1162/neco_a_01357.
- Cansu Sancaktar, Marcel van Gerven, and Pablo Lanillos. End-to-end pixel-based deep active inference for body perception and action, 2020.
- Philipp Schwartenbeck, Johannes Passecker, Tobias U Hauser, Thomas H B FitzGerald, Martin Kronbichler, and Karl Friston. Computational mechanisms of curiosity and goal-directed exploration. *bioRxiv*, 2018. doi: 10.1101/411272. URL <https://www.biorxiv.org/content/early/2018/09/07/411272>.
- Peter Uhlhaas, Gordon Pipa, Bruss Lima, Lucia Melloni, Sergio Neuenschwander, Danko Nikolić, and Wolf

Singer. Neural synchrony in cortical networks: history, concept and current status. *Frontiers in integrative neuroscience*, 3:17, 2009.

Jonathan S. Yedidia. Message-passing algorithms for inference and optimization. *Journal of Statistical Physics*, 145(4):860–890, Nov 2011. ISSN 1572-9613. doi: 10.1007/s10955-011-0384-7. URL <https://doi.org/10.1007/s10955-011-0384-7>.

John Winn and Christopher Bishop. Variational message passing. *Journal of Machine Learning Research*, 6:661–694, 2005.

Ozan Çatal, Tim Verbelen, Johannes Nauta, Cedric De Boom, and Bart Dhoedt. Learning perception and planning with deep active inference, 2020.

Appendix A: $BTAI_{3MF}$ algorithm

Algorithm 1: $BTAI_{3MF}$: action-perception cycles (with relevant equations indicated in round brackets).

Input: env the environment,

$O_0 = \{O_0^o \mid o = 1, \dots, \#O\}$ the initial observations,

$\mathbf{A} = \{\mathbf{A}^o \mid o = 1, \dots, \#O\}$ the likelihood mapping of each observation,

$\mathbf{B} = \{\mathbf{B}^s \mid s = 1, \dots, \#S\}$ the transition mapping for each hidden state,

$\mathbf{C} = \{\mathbf{C}^i \mid i = 1, \dots, N\}$ the prior preferences of each subset of observations,

$\mathbf{D}_0 = \{\mathbf{D}_0^s \mid s = 1, \dots, \#S\}$ the prior over each initial state,

N the number of planning iterations,

M the number of action-perception cycles.

```

 $P(S_0^s|O_0) \leftarrow \text{I-step}(O_0, \mathbf{A}, \mathbf{D}_0)$  // I-step from Section 2.3.1
root  $\leftarrow$  CreateTreeNode(
    beliefs =  $P(S_0^s|O_0)$ , action = -1, cost = 0, visits = 1
) // Create the root node for the MCTS, where -1 is a dummy value
repeat  $M$  times
    repeat  $N$  times
        node  $\leftarrow$  SelectNode(root) // Using (2) recursively
        eNodes  $\leftarrow$  ExpandChildren(node,  $\mathbf{B}$ ) // P-step from Section 2.3.2 for each action
        Evaluate(eNodes,  $\mathbf{A}, \mathbf{C}$ ) // Compute (1) for each expanded node
        Backpropagate(eNodes) // Using (3) and (4)
    end
     $\mathbf{a}^* \leftarrow$  SelectAction(root) // Using (5)
     $O_{t+1} \leftarrow env.Execute(\mathbf{a}^*)$ 
    child  $\leftarrow root.children[\mathbf{a}^*]$  // Get root child corresponding to  $\mathbf{a}^*$ 
     $P(S_{t+1}^s) \leftarrow child.beliefs$  // Get the empirical prior  $P(S_{t+1}^s) = \text{Cat}(\mathbf{D}_{t+1}^s)$ 
     $P(S_{t+1}^s|O_{t+1}) \leftarrow \text{I-step}(O_{t+1}, \mathbf{A}, \mathbf{D}_{t+1})$  // I-step from Section 2.3.1
    root  $\leftarrow$  CreateTreeNode(
        beliefs =  $P(S_{t+1}^s|O_{t+1})$ , action =  $\mathbf{a}^*$ , cost = 0, visits = 1
    ) // Create the root node of the next action-perception cycle
end

```
