
Optimal Selective Attention in Reactive Agents

Technical Report

Roy Fox

Naftali Tishby

School of Computer Science and Engineering
The Hebrew University

Abstract

In POMDPs, information about the hidden state, delivered through observations, is both valuable to the agent, allowing it to base its actions on better informed internal states, and a "curse", exploding the size and diversity of the internal state space. One attempt to deal with this is to focus on reactive policies, that only base their actions on the most recent observation. However, even reactive policies can be demanding on resources, and agents need to pay selective attention to only some of the information available to them in observations. In this report we present the minimum-information principle for selective attention in reactive agents. We further motivate this approach by reducing the general problem of optimal control in POMDPs, to reactive control with complex observations. Lastly, we explore a newly discovered phenomenon of this optimization process — period doubling bifurcations. This necessitates periodic policies, and raises many more questions regarding stability, periodicity and chaos in optimal control.

1 Introduction

For an intelligent agent interacting with its environment, information is valuable. By observing and retaining information about its environment, the agent can form beliefs and make predictions. It represents these beliefs in an internal state, on which it can then base its actions.

If information about some event in the world is unavailable to the agent, through the lack of observability or attention, its internal state is independent of that event, and so are its actions, potentially incurring otherwise avoidable costs. The same is true if the

information is only partially available, limiting the extent to which the agent's actions can depend on the state of the world.

However, information is also a "curse". Retaining much information about the world requires the agent to have a large and rich internal state space, representing diverse beliefs. This leads to complex policies for inference and control, which are computationally hard both to find and to apply. Designed agents should not be — and evolved agents are unlikely to be — more complex than is sufficient for them to perform well.

The "curse of dimensionality" [3] is the challenge of representing in the internal state space the entire belief space — the space of probability distributions over world states. The volume of this simplex is exponential in the number of world states, and approximate methods [19] [17] [1] [12] are required to explore and represent policies over this space.

The "curse of history" [14] results from representing only reachable Bayesian beliefs — posteriors of the world state given each possible observable history. The Bayesian belief is a sufficient statistic of the observable history for the world state, keeping all available information about it. Unfortunately, the size of this space can be exponential in the length of the history.

This realization immediately suggests the idea of truncating the observable history by forgetting older observations. Taken to the extreme, this leads to reactive agents [10] [20], in which each internal state can only take into account the most recent observation, discarding the previous internal state. The internal state space of reactive agents needs not be larger than the observation space, which removes the curse of history in domains where the set of observations is not too large.

Definition 1. *A reactive agent bases its actions only on the most recent observation. In contrast, a retentive agent can base its actions on a memory state, which is updated with each observation, and thus sum-*

marizes the entire observable history.

A drawback of this approach is that, since the history is no longer grounded in a known initial belief, a new challenge arises of identifying which beliefs these internal states represent. This challenge generally requires forward-backward algorithms [6], as opposed to fully observable Markov Decision Processes which are solvable by backward (dynamic programming) algorithms [3].

In addition, the original difficulty remains in domains where the observation space is still too large, such as the one presented in Section 3. In this sense, the curse of history is a special case of the following principle, which we might call the "curse of information".

An agent's input — its sensors, and its memory when available — usually contains too much information for the agent to process. For the agent to encode all of this information in its new internal state, an internal state space is required that is too large to be manageable and utilized by feasible policies. As a matter of practicality, an agent must have selective attention. A retentive agent must also have selective retention [6], which is beyond the scope of this paper.

Definition 2. *A reactive agent (similarly, a retentive agent), is said to have selective attention (resp. selective retention) if its internal state has less information about the world state than its observation (resp. observable history) does.*

Reactive policies have been explored before in [10], with some of their challenges noted in [20]. A policy-gradient algorithm for finding such policies was presented in [7], which has the nice property of avoiding the forward-backward coupling. However, the local optimum it finds is not guaranteed to be a fixed point of the value recursion.

Information considerations in dynamical systems were presented in [24]. Algorithms were later introduced for trading off value and information in fully observable Markov Decision Processes [18] and in partially observable ones where actions have no external effect [6].

This paper offers three novel contributions, in each of the following sections.

Section 3 shows that reactive policies are as expressive as retentive policies, under proper redefinition of the model. This motivates our focus on reactive agents, at the same time that it demands a more principled cure for the curse of information than simply discarding the memory.

Section 4 provides such a principle, namely the minimum-information principle. We present the prin-




Figure 1: Schematic model of a reactive agent interacting with its environment

ciple and formalize it, discuss its relation to source coding, and give an algorithm for its numeric solution.

Section 5 demonstrates a newly discovered phenomenon in optimal control, namely the occurrence of bifurcations when attention is traded off with external cost. This is illustrated using two examples.

We conclude with a short discussion of these contributions and their consequences.

2 Preliminaries

We model the interaction of an intelligent agent with its environment using the formalism of Partially Observable Markov Decision Processes (POMDPs). A POMDP is a discrete-time dynamical system with state $s_t \in \mathcal{S}$. In time t , the system emits an observation $o_t \in \mathcal{O}$ with probability $\sigma(o_t|s_t)$. It then receives from the interacting agent an input action $a_t \in \mathcal{A}$, and transitions to a new state s_{t+1} with probability $p(s_{t+1}|s_t, a_t)$. For our purposes here, the sets \mathcal{S} , \mathcal{O} and \mathcal{A} are finite, and we are only concerned with stationary (time-invariant) POMDPs, where the model parameters p and σ are fixed for every time step.

A reactive agent has no internal memory state, and can only base its actions on the most recent observation. The agent consists of two modules, the sensor making the observation o_t and the actuator taking the action a_t (Figure 1). The reactive policy π of the agent is implemented by linking the two modules through a communication channel, such that the action a_t is taken with probability $\pi_t(a_t|o_t)$ in reaction to observation o_t . The policy is called periodic with period \mathcal{T} if $\pi_t = \pi_{t+\mathcal{T}}$ for every time step t . The policy is called stationary if it has period 1, i.e. π_t is fixed for every time step.

The model and the policy together induce a stochastic process over the variables $\{s_t, o_t, a_t\}$ (Figure 2). Due to the agent's lack of memory, the states $\{s_t\}$ form a Markov chain. In the following we always assume that the process is ergodic. This implies that, if the agent policy has period \mathcal{T} , then for each phase $0 \leq t < \mathcal{T}$




Figure 2: Graphical model of a reactive agent interacting with its environment

there exists a unique marginal distribution $\bar{p}_t(s_t)$ that is a fixed point of the \mathcal{T} -step forward recursion

$$\bar{p}_t(s_{t+\mathcal{T}}) = \sum_{s_t} \bar{p}_t(s_t) P_{t,\pi}(s_{t+\mathcal{T}}|s_t)$$

with

$$P_{t,\pi}(s_{t+\mathcal{T}}|s_t) = \sum_{s_{t+1}, \dots, s_{t+\mathcal{T}-1}} \prod_{\tau=t}^{t+\mathcal{T}-1} P_{\pi_\tau}(s_{\tau+1}|s_\tau)$$

and

$$P_{\pi_\tau}(s_{\tau+1}|s_\tau) = \sum_{o_\tau, a_\tau} \sigma(o_\tau|s_\tau) \pi_\tau(a_\tau|o_\tau) p(s_{\tau+1}|s_\tau, a_\tau).$$

These marginal distributions are therefore periodic with the same period \mathcal{T} , i.e. $\bar{p}_t = \bar{p}_{t+\mathcal{T}}$, and inside a cycle the phases are linked through the 1-step forward recursion

$$\bar{p}_{t+1}(s_{t+1}) = \sum_{s_t} \bar{p}_t(s_t) P_{\pi_t}(s_{t+1}|s_t). \quad (1)$$

The marginal distributions also induce beliefs

$$b_t(s_t|o_t) = \frac{\bar{p}_t(s_t) \sigma(o_t|s_t)}{\bar{\sigma}_t(o_t)},$$

with

$$\bar{\sigma}_t(o_t) = \sum_{s_t} \bar{p}_t(s_t) \sigma(o_t|s_t).$$

The belief is the posterior distribution of the state given the observation.

In this paper we will have the agent incur an external nominal cost $c(s_t, a_t)$ when it takes action a_t in state s_t , and measure the quality of a policy by the long-term average expected cost

$$\mathcal{C} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[c(s_t, a_t)]$$

in the stochastic process that the policy induces. If the policy has period \mathcal{T} and the process is at its periodic marginal distribution, then

$$\mathcal{C} = \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} \sum_{s_t, o_t, a_t} \bar{p}_t(s_t) \sigma(o_t|s_t) \pi_t(a_t|o_t) c(s_t, a_t).$$




Figure 3: Graphical model of a retentive agent interacting with its environment




Figure 4: Reduction from a retentive policy to a reactive policy

This undiscounted expected cost is appropriate for studying stationary processes. In contrast, discounting the cost by γ^t emphasizes transient effects, up to horizon $O(\frac{1}{1-\gamma})$. A related fault with discounting in reactive policies is discussed in [20].

3 Reduction from retentive to reactive policies

Consider a retentive agent [2] [6] interacting with a POMDP (Figure 3). The agent has an internal state $m_t \in \mathcal{M}$, and an inference policy q_t controlling it, such that with probability $q_t(m_t|m_{t-1}, o_t)$ the memory state m_{t-1} is updated to m_t upon observing o_t in time step t . The control policy $\pi_t(a_t|m_t)$ is allowed to depend not only on the most recent observation, but on the summary of the entire observable history represented in m_t .

In a given POMDP, retentive policies (q, π) are much more expressive and powerful than reactive policies. Interestingly, however, there exists another (time-variant) POMDP in which $\pi' = (q, \pi)$ can be implemented as a reactive policy (Figure 4). This new POMDP is similar in spirit to the cross-product MDP [11], and the distinction between them is discussed below.

Formally, let the state space of the new POMDP be $\mathcal{S}' = \mathcal{M} \times \mathcal{S}$, the observation space $\mathcal{O}' = \mathcal{M} \times (\mathcal{O} \cup \{\perp\})$, where \perp is some null-observation symbol, and the action space $\mathcal{A}' = \mathcal{M} \cup \mathcal{A}$. Let the dynamics advance at twice the frequency, with each time step tak-

ing half as long. The state at time t is $s'_t = (m_{t-1}, s_t)$, and it emits an observation with distribution

$$\sigma'_t((m_{t-1}, o_t)|(m_{t-1}, s_t)) = \sigma(o_t|s_t).$$

The agent, upon observing (m_{t-1}, o_t) , can apply its inference policy to generate the next memory state m_t . It then takes the "action" of committing m_t to "external storage"

$$p'_t((m_t, s_t)|(m_{t-1}, s_t), m_t) = 1.$$

In this new state at time $t + \frac{1}{2}$, the committed memory state is observable

$$\sigma'_{t+\frac{1}{2}}((m_t, \perp)|(m_t, s_t)) = 1,$$

and the agent can apply its control policy to take the action a_t , thus controlling the transition

$$p'_{t+\frac{1}{2}}((m_t, s_{t+1})|(m_t, s_t), a_t) = p(s_{t+1}|s_t, a_t).$$

Note that it should be inadmissible for the agent to commit a memory state in a non-integer time step, or take an action in an integer one. This can be enforced by penalizing the wrong type of action, which is the main reason that the new POMDP needs to be time-variant.

Assuming that the agent follows these restrictions, the new POMDP induces the same stochastic process over the variables $\{s_t, o_t, m_t, a_t\}$ as the original one for any given policy, establishing the reduction.

Our reduction is related to the cross-product MDP of [11]. However, the two models have different formulations that serve their different purposes — where the cross-product MDP creates structure to be exploited in planning algorithms, our formulation flattens this structure to reduce the problem to a simpler one. To achieve this, instead of the implicit restriction in [11] that policies depend only on the agent state, we model the same constraint explicitly as partial observability. Furthermore, by breaking each time step into two we avoid the exponential action space of the cross-product MDP.

Lastly, an important issue to consider is the memory state space \mathcal{M} . The standard approach in the reinforcement learning literature is to have \mathcal{M} be the belief space, the simplex of distributions over \mathcal{S} , and q the Bayesian inference policy¹. Such a choice would make \mathcal{S}' , \mathcal{O}' and \mathcal{A}' uncountable, as opposed to our usual assumption that these sets are finite.

¹It is also common to have actions as part of the observable history, which our notation allows but does not require.

Alternatively, we can have in \mathcal{M} only reachable beliefs. If the support of the inference policy q is finite², then over a finite horizon only a finite number of beliefs are reachable. Unfortunately, due to the "curse of history", this number is exponential in the horizon, which renders this reduction — and indeed many existing approaches to POMDPs — impractical.

This difficulty underlines the need for selective attention. Theoretically, the support of m_t needs never be larger than that of (m_{t-1}, o_t) , at least in terms of sufficient inference. However, it should practically be much smaller than that — roughly the same size as the support of m_{t-1} — if the agent is to interact with the system for significant horizons without exploding in complexity. The ability of the agent to selectively attend to its input, whether from sensors or from memory, and to retain not all, but only the most useful information, is key to reducing this complexity.

This is the approach taken by Finite State Controllers (FSCs) [15], where the number of memory states is fixed. Several heuristic algorithms exist for finding a good FSC, however this problem is hard and highly non-convex. The policy of a FSC is time-invariant, and as we see in Section 5 a stationary Bellman-optimal solution is generally not stable.

4 Minimum-information principle

Our guiding principle in formalizing selective attention is the reduction of information complexity, as measured by the Shannon mutual information between the observation o_t and the action a_t . We first present the principle, and then justify it by relating it to source coding. We note that numerous other justifications and connections exist, some discussed previously [9] [24] [25] [16] [8], and some should be explored further, particularly in the context of POMDP planning.

The pointwise mutual information between o_t and a_t in time step t is given by

$$i_t(o_t, a_t) = \log \frac{\pi_t(a_t|o_t)}{\bar{\pi}_t(a_t)},$$

with

$$\bar{\pi}_t(a_t) = \sum_{o_t} \bar{\sigma}_t(o_t) \pi_t(a_t|o_t). \quad (2)$$

This can be thought of as the internal informational cost of choosing action a_t in reaction to observation o_t . The long-term average expectation of this internal cost, similar to the external cost, is

$$\mathcal{I} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[i_t(o_t, a_t)].$$

²For example, the Bayesian inference policy is deterministic.

If the policy has period \mathcal{T} and the process is at its periodic marginal distribution then

$$\begin{aligned}\mathcal{I} &= \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} \sum_{o_t, a_t} \bar{\sigma}_t(o_t) \pi_t(a_t|o_t) i_t(o_t, a_t) \\ &= \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} D_{\text{KL}}[\pi_t || \bar{\pi}_t] = \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} \mathbb{I}[o_t; a_t].\end{aligned}$$

Here $D_{\text{KL}}[\pi_t || \bar{\pi}_t]$ is the Kullback-Leibler divergence of π_t from $\bar{\pi}_t$, and $\mathbb{I}[o_t; a_t]$ is the Shannon mutual information between o_t and a_t .

$D_{\text{KL}}[\pi_t || \bar{\pi}_t]$ is a measure of the cognitive effort required for the agent to diverge from a passive, uncontrolled policy $\bar{\pi}_t$ to an active, controlled policy π_t . Unlike [9] [24] [25], we allow the passive policy, as well as the active one, to be designed or evolved. The uncontrolled policy that minimizes the informational cost \mathcal{I} is the appropriate marginal distribution of the action (2) [5].

Among agents incurring external cost $\mathcal{C} \leq C$, the simplest agent, in some sense, minimizes the internal cost \mathcal{I} . In other words, the agent needs to trade off its external and internal costs. To link these views, the Lagrange multiplier β corresponding to the constraint $\mathcal{C} \leq C$ in the optimization of \mathcal{I} is a conversion rate between the two types of cost. We can then write the total cost as

$$\mathcal{F} = \frac{1}{\beta} \mathcal{I} + \mathcal{C}.$$

\mathcal{F} is called the free energy, due to its similarity to the quantity of the same name in statistical physics, with β taking the part of the inverse temperature.

For a given β , the agent chooses its policy so as to minimize the free energy, under two constraints. First, the dynamics of the system follow the forward recursion (1). Second, \bar{p}_t , $\pi_t(\cdot|o_t)$ and $\bar{\pi}_t$ need to be probability distributions, each summing to 1. The constraints that they are non-negative can be ignored, since they will be either inactive or weakly active.

This gives for horizon T the Lagrangian $\mathcal{L}_{\bar{p}, \pi, \bar{\pi}}$

$$\begin{aligned}&= \frac{1}{T} \sum_{t=0}^{T-1} \left(\sum_{s_t, o_t, a_t} \bar{p}_t(s_t) \sigma(o_t|s_t) \pi_t(a_t|o_t) f_t(s_t, o_t, a_t) \right. \\ &+ \sum_{s_{t+1}} \nu_{t+1}(s_{t+1}) \left(\sum_{s_t} \bar{p}_t(s_t) P_{\pi_t}(s_{t+1}|s_t) - \bar{p}_{t+1}(s_{t+1}) \right) \\ &- \varphi_t \left(\sum_{s_t} \bar{p}_t(s_t) - 1 \right) + \eta_t \left(\sum_{a_t} \bar{\pi}_t(a_t) - 1 \right) \\ &\left. + \sum_{o_t} \lambda_t(o_t) \left(\sum_{a_t} \pi_t(a_t|o_t) - 1 \right) \right),\end{aligned}$$

with

$$f_t(s_t, o_t, a_t) = \frac{1}{\beta} i_t(o_t, a_t) + c(s_t, a_t).$$

4.1 Necessary conditions for optimality

This optimization problem is far from convex, and no efficient algorithm is known for finding the global optimum. Indeed, as β tends to infinity, the agent's policy becomes deterministic, and some problems involving deterministic reactive policies are known to be NP-complete [10].

Nevertheless, we can consider local minima by finding the first-order necessary conditions for a solution to be optimal. That is, we differentiate the Lagrangian by each of its parameters, and require that this derivative equals 0.

For \bar{p} , this gives us a backward recursion

$$\begin{aligned}\nu_t(s_t) &= \sum_{o_t, a_t} \sigma(o_t|s_t) \pi_t(a_t|o_t) f_t(s_t, o_t, a_t) \\ &+ \sum_{s_{t+1}} P_{\pi_t}(s_{t+1}|s_t) \nu_{t+1}(s_{t+1}) - \varphi_t.\end{aligned}\quad (3)$$

Due to overconstraining, we have some degrees of freedom in choosing the multipliers to satisfy the Karush-Kuhn-Tucker conditions [4]. If the policy has period \mathcal{T} , we will choose φ_t to also have period \mathcal{T} and satisfy

$$\frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} \varphi_t = \mathcal{F},$$

so that ν_t also has period \mathcal{T} . Thus $\nu_t(s_t)$ measures the fluctuation from the average free energy \mathcal{F} of the state s_t in phase t of the cycle.

The first-order necessary conditions for π are

$$\pi_t(a_t|o_t) = \frac{1}{Z_t(o_t)} \bar{\pi}_t(a_t) \exp(-\beta d_t(o_t, a_t)) \quad (4)$$

with

$$\begin{aligned}d_t(o_t, a_t) &= \sum_{s_t} b_t(s_t|o_t) c(s_t, a_t) \\ &+ \sum_{s_t, s_{t+1}} b_t(s_t|o_t) p(s_{t+1}|s_t, a_t) \nu_{t+1}(s_{t+1})\end{aligned}$$

and the normalizing partition function

$$Z_t(o_t) = \sum_{a_t} \bar{\pi}_t(a_t) \exp(-\beta d_t(o_t, a_t)),$$

and for $\bar{\pi}$ we have (2) as promised.

As β tends to infinity, the optimal policy in (4) becomes deterministic. Together with (3), it becomes a Bellman equation [3].

For finite β , on the other hand, the optimal policy is stochastic, which is a welcome outcome in many respects. The best deterministic reactive policy is generally arbitrarily worse than the optimal stochastic




Figure 5: Reactive policy as a source-coding problem

reactive policy [20]. Optimality in reactive policies requires stochasticity. Unfortunately, many planning algorithms rely on the smaller space of deterministic policies (e.g. [14]), and others lack a principle by which to gauge the optimal amount of uncertainty in the agent’s actions (e.g. [2]). We propose minimum information as such a principle.

Furthermore, in practice the model used for planning is itself uncertain. Using deterministic policies could overfit to the available data, and hinder further learning [22]. Information considerations provide a principled way of fitting the uncertainty of the policy to the uncertainty of the model [18].

In reinforcement learning, it is common to use softmax to obtain stochastic planning and exploration policies [2] [22]. Note that softmax is a special case of (4), with the uniform prior instead of the marginal $\bar{\pi}_t$. That information theory provides a better principle for stochasticity in optimization is further illustrated in the next subsection.

4.2 Sequential rate-distortion

The form of (4) and (2) may be familiar as the solution to the rate-distortion problem of lossy source coding [5]. Indeed, minimum-information optimal control can be construed as a sequential rate-distortion problem [23].

The reactive agent’s policy is a channel from its sensor to its actuator (Figure 5). Following the classic model of such a channel, the sensor can be considered an encoder which, upon observing the "source" o_t , chooses a "codeword" m_t . It transmits it to the actuator, a decoder which then "reconstructs" the intended a_t .

For a given time step t , and with a source distribution $\bar{\sigma}_t(o_t)$ and a distortion function $d_t(o_t, a_t)$ fixed, this would be a standard source-coding problem. Let a feasible agent be one achieving at most D expected distortion

$$\sum_{o_t, a_t} \bar{\sigma}_t(o_t) \pi_t(a_t | o_t) d_t(o_t, a_t) \leq D.$$

Now suppose we are interested in the feasible agent with the simplest internal state space, as measured by the size of the "codebook" \mathcal{M} . The rate-distortion theorem [5] states that the simplest feasible agent is the one minimizing $I[o_t; a_t]$.

In a sequential rate-distortion problem, the solution π_t in time step t affects future source distributions $\bar{\sigma}_\tau$ in time steps $\tau > t$, as well as past distortions d_τ in time steps $\tau < t$. This creates a coupling between the forward inference process of computing marginal distributions, and the backward control process of computing value functions. This is further complicated in partially observable processes, where d_t also depends on the forward inference process through the beliefs b_t .

Coupling makes sequential rate-distortion complex, both conceptually and computationally. Conceptually, the results of rate-distortion theory are no longer known to hold in the sequential case. If we nevertheless accept the minimum information as a solid guiding principle in our optimization, we find that this optimization is computationally hard. We can optimize the policy in each time step given the other time steps with algorithms like Blahut-Arimoto [13]. However the forward-backward algorithm for finding the overall policy is only guaranteed to converge to a local optimum [6].

4.3 Optimization algorithm

The forward recursion (1), the backward recursion (3), the optimal policy (4) and its marginal (2) are necessary conditions for a solution to be optimal. They also provide an algorithm for finding a good solution: iteratively plug the current solution in the right-hand side of one of the equations, to obtain a better solution, until (asymptotically) no such improvement is possible. Many existing algorithms employ a similar scheme. For example, in the Generalized Policy Iteration algorithm for planning in MDPs [22], there is some schedule for alternating between³ policy evaluation, a variant of (3), and policy improvement, a variant of (4) with $\beta \rightarrow \infty$.

A sophisticated schedule can guarantee that the solution improves monotonically with each iteration [6]. Here we suggest the following simpler schedule, for which such a guarantee does not hold, but which empirically converges to good solutions in practice.

Repeat until convergence:

1. Compute the marginal $\bar{\pi}$ given the current solution for π , by applying (2).
2. Compute the value function ν given the current solution for \bar{p} , π and $\bar{\pi}$. This can be done by iteratively applying (3) until it converges, or by solving it as a system of linear equations.

³A forward equation is not needed in fully observable problems if attention is not selective.

3. In a forward algorithm, until convergence to a limit cycle:
 - (a) Compute the marginal \bar{p}_t given the current solution for \bar{p}_{t-1} and π_{t-1} , by applying (1).
 - (b) Compute the optimal policy π_t given the current solution for \bar{p}_t , $\bar{\pi}$ and ν , by applying (4).

5 Periodicity in reactive policies

Throughout the previous sections, we always referred to periodic reactive policies rather than stationary ones, even though the POMDP itself is assumed to be stationary. Periodic reactive policies may seem to be a contradiction in terms, since their actions depend not only on the most recent observation, but also on the time t . They require a clock to be available to the actuator, with period that is a multiple of the policy period.

We argue that periodic policies must inevitably be a part of the solution concept of POMDPs with selective attention. When paying full attention to inputs, in the form of exact Bayesian inference, we can restrict the discussion to stationary policies [19]. When attention is partial, there are significant drawbacks to considering only stationary policies.

One drawback is that the best stationary policy is generally arbitrarily worse than the optimal periodic policy. Adapting the example in [20], consider the POMDP illustrated in Figure 6. This model has 2 states, 1 (uninformative) observation and 2 actions. The actions deterministically set the next state, and a reward (negative cost) is given for switching to the other state.

The optimal stationary retentive policy for this POMDP is to have two internal memory states, each indicating a different action, and switch between them in each time step. This policy gets the reward in each time step, but incurs 1 bit of internal cost⁴.

On the other hand, a stationary reactive policy in an unobservable POMDP is just a fixed distribution over the actions, and it can be no better in this instance than the uniform distribution. This policy yields only half the expected reward, but incurs no internal cost.

Lastly, the reactive policy of period 2 which alternates between the actions also receives the full reward, at seemingly no internal cost. In fact, this would seemingly also be the preferred retentive solution, if the internal cost is taken into consideration.

Of course, counting no internal cost for a periodic pol-

⁴See [6] for the definition of the internal cost of a retentive policy.

icy is cheating. Instead of paying attention to its sensors or memory, the agent is paying attention to a clock, but that attention is still a burden on internal resources.

Similar to the informational cost between o_t and a_t , we need to add a term for the informational cost between t and a_t . For a reactive policy with period \mathcal{T} , this cost term can naturally be defined by

$$\begin{aligned} I[t; a_t] &= \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} D_{\text{KL}}[\bar{\pi}_t \| \bar{\pi}] \\ &= \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} \sum_{a_t} \bar{\pi}_t(a_t) \log \frac{\bar{\pi}_t(a_t)}{\bar{\pi}(a_t)}, \end{aligned}$$

with

$$\bar{\pi}(a) = \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} \bar{\pi}_t(a).$$

Here we use the fact that the phase of the cycle is distributed uniformly during the process.

Adding the term $I[t; a_t]$ to the free energy is equivalent to asserting that a clock is observable to the agent, and that attention to it is as costly as to any other part of the observation. The pointwise informational cost is now

$$\tilde{i}_t(o_t, a_t) = \log \frac{\pi_t(a_t|o_t)}{\bar{\pi}(a_t)},$$

and the average expected internal cost is

$$\begin{aligned} \tilde{I} &= \frac{1}{\mathcal{T}} \sum_{t=0}^{\mathcal{T}-1} I[o_t; a_t] + I[t; a_t] \\ &= I[o_t; a_t|t] + I[t; a_t] = I[t, o_t; a_t]. \end{aligned}$$

The values of \tilde{f}_t , $\tilde{\nu}_t$ and \tilde{d}_t change accordingly, and the optimal policy is now

$$\pi_t(a_t|o_t) = \frac{1}{\tilde{Z}_t(o_t)} \bar{\pi}(a_t) \exp(-\beta \tilde{d}_t(o_t, a_t)),$$

with the proper partition function $\tilde{Z}_t(o_t)$.

This allows us to consider policies which are "softly periodic", in that they attend to some but not all time information. Figure 7 shows the information-cost curve for the POMDP in Figure 6, and Figure 8 shows the final-state diagram for the iterative algorithm with the schedule in Section 4.3.

Interestingly, this problem exhibits a bifurcation at $\beta = 1$. Below this value, information is too costly, and the optimal solution is the stationary uniform policy. At $\beta = 1$, the system undergoes a period-doubling bifurcation, and above this value the optimal policy becomes periodic with period 2 — the two phases of




Figure 6: An unobservable POMDP, having an optimal policy of period 2




Figure 7: Information-cost curve for the POMDP in Figure 6; Points on the curve were achieved by the algorithm with different values of β , and points above it are achievable; The curve is convex, with slope $-\beta$

the cycle are given by the two branches in Figure 8. The "hardness" of this periodicity, as measured by the information $I[t; a_t]$, grows continuously from 0, and tends to 1 bit as β tends to infinity.

Above the critical point, a third solution exists, which is a fixed point of the optimization schedule (the dashed line in Figure 8). This solution is the optimal stationary reactive policy, but it is an unstable fixed point: starting the optimization from a small perturbation of this solution does not converge back to it, but diverges until it reaches the periodic solution. Thus we have a supercritical pitchfork bifurcation [21].

The instability of the stationary solution is another paramount reason for allowing periodic policies. It would be practically impossible to find a stationary solution using Bellman-like variational methods, as the one presented in this paper. In contrast, approaches such as policy-gradient methods [2] generally can find stationary solutions, but these are generally not fixed points of a Bellman recursion, and are thus not Bellman-optimal [3].

5.1 Robot example

As another example, consider the POMDP illustrated in Figure 9. Here a robot is engaged in moving items




Figure 8: Finite-state diagram for the iterative algorithm applied to the POMDP in Figure 6, as a function of the cost conversion rate β ; Points on the curve are the probability of taking the action "right" in each phase of the limit cycle of the algorithm, when run to convergence with the given β




Figure 9: A POMDP of a robot moving items from the left end of a corridor to the right one; Shown actions succeed with probability 0.8, otherwise the state remains the same; Location sensor correct with probability 0.88, load sensor with probability 0.7




Figure 10: Information-cost curve for the POMDP in Figure 9




Figure 11: Finite-state diagram for the POMDP in Figure 9; Points on the curve are the probability of taking the action "unload" when the sensors indicate "right end" and "loaded", in each phase of the limit cycle; The lower branch is grayed out for clarity

from the left end of a corridor to the right one. The robot can be in one of 4 states: it can be at either end of the corridor, and it can be carrying an item or not. It has 4 actions: to move to the left end of the corridor, or to the right, or to pick up or put down an item. However, an action can fail with probability 0.2, leaving the robot at the same state. The robot can only pick up an item at the left end of the corridor, and it receives a reward for dropping an item at the right end.

The robot has 4 possible observations from two binary sensors, telling it its position and whether or not it is carrying an item. The location sensor is more reliable, showing the correct position with probability 0.88. The load sensor only shows the correct load state with probability 0.7. The parameters were selected for visual clarity of the results.

Figure 10 shows the information-cost curve for this problem. Here there are two phase transitions, where the period doubles to 2, then again to 4 (Figure 11). When attention is scarce, the robot's actions are more uniformly random. In this situation, the sensors, although noisy, carry more relevant information than the clock, since they better correspond to the actual state.

As attention increases with β , the robot relies more and more on its sensors. Conditional on the observation, the robot makes its actions less and less stochastic. At some point, the policy is reliable enough that the clock has more relevant information about the load state than the noisy sensor. At that point, a pitchfork bifurcation occurs, and the robot begins to rely mostly on the parity of the clock to decide when to move, and

on the location sensor to decide where to move and whether to load or unload.

With the parameters above, as β keeps increasing, the clock eventually becomes even more reliable than the location sensor, and a second period doubling occurs, to period 4. Asymptotically as β grows to infinity, the clock signal takes precedence over both sensors, and the agent unloads its cargo on schedule even if its sensors tell it that it is dislocated or empty handed.

6 Discussion

In this paper we presented three novel results involving reactive agents interacting with partially observable systems. We have motivated the focus on reactive policies through a reduction from retentive policies, introduced a principle and an algorithm for optimizing reactive policies, and explored a surprising aspect of their phenomenology.

We conclude with a few remarks on the implication of each contribution.

6.1 Selective attention as clustering

Information-constrained clustering can also be construed as source coding [16], so that the data to be clustered is considered the source, and the cluster centroids the reconstruction. Following the relation we show between selective attention and source coding, we can think of a reactive policy as a soft clustering of observations into actions.

With the information constraint removed, the clustering becomes hard, mapping each data point to its closest centroid. Similarly in our case, as β grows the policy becomes more deterministic, until at $\beta \rightarrow \infty$ it always picks the optimal action for each observation.

The implication of viewing reactive policies as clustering is that actions should generally be simpler, and never more complex, than the observations on which they are based. Indeed, there is a duality between observations and actions, and between selective attention (the retained part of the observation) and selective action (the intended part of the action, as divergence from the prior $\bar{\pi}$). Information that is not retained cannot be used for choosing actions, and there is no point in retaining information that is not used.

6.2 Implications of selective attention for retentive agents

In this paper we have focused on reactive agents, and introduced the minimum-information principle for optimal selective attention. However, as the reduction

in Section 3 shows, this has implications for retentive agents as well.

The effect of selective attention is to make internal states less complex than their inputs, by discarding information that is not useful enough. When applied to the inference policy, this leads to approximate inference, that trades off the external value of information in guiding actions with its internal cost in information complexity. In fact, an inference process in POMDPs is equivalent to sequential clustering. With each new observation o_t , the pair (m_{t-1}, o_t) is clustered into a new internal state m_t .

The major challenge when planning in POMDPs is approximating the Bayesian belief in such a way that allows efficient planning and execution, while not losing too much value. Selective attention, and in this case retention, is precisely such a principle. The application of this approach to retentive agents is left for future work.

6.3 Policy bifurcations and chaos theory

We have discovered the occurrence of bifurcations in the optimization process of reactive policies. It presents many of the characterizing features of chaos theory of iterated functions, such as period doubling and slow convergence near the bifurcation points. We expect to see many more such features in other, more complex systems. We conjecture that systems with more states, perhaps infinitely many, can present a cascade of bifurcations, leading to aperiodicity and chaos.

A full investigation of the bearings of the theories of bifurcation and chaos to optimal control in dynamical systems is beyond the scope of this report. To the extent that such a connection exists, it could be of profound philosophical implications, as it could indicate that intelligent agents interacting with complex environments must choose among the following alternatives:

- Plan with very little attention of their inputs
- Plan for very short horizons
- Plan with some degree of inability to identify their own value function or predict their own future actions.

References

- [1] Douglas Aberdeen. A (revised) survey of approximate methods for solving partially observable Markov decision processes. *National ICT Australia, Canberra, Australia*, 2003.
- [2] Douglas A. Aberdeen. *Policy-gradient algorithms for partially observable Markov decision processes*. PhD thesis, Australian National University, 2003.
- [3] Richard Bellman. *Dynamic programming*. Princeton University Press, 1957.
- [4] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [5] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley Series in Telecommunications and Signal Processing. John Wiley & Sons, 2012.
- [6] Roy Fox and Naftali Tishby. Bounded planning in passive POMDPs. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1775–1782, 2012.
- [7] Tommi Jaakkola, Satinder P. Singh, and Michael I. Jordan. Reinforcement learning algorithm for partially observable Markov decision problems. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, pages 345–352, 1995.
- [8] Edwin T. Jaynes. *Probability theory: The logic of science*. Cambridge University Press, 2003.
- [9] Hilbert J. Kappen, Vicenç Gómez, and Manfred Opper. Optimal control as a graphical model inference problem. *Machine Learning*, 87(2):159–182, 2012.
- [10] Michael L. Littman. Memoryless policies: Theoretical limitations and practical results. In *From Animals to Animats 3: Proceedings of the third international conference on simulation of adaptive behavior (SAB)*, volume 3, page 238, 1994.
- [11] Nicolas Meuleau, Kee-Eung Kim, Leslie Pack Kaelbling, and Anthony R Cassandra. Solving pomdps by searching the space of finite policies. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pages 417–426. Morgan Kaufmann Publishers Inc., 1999.
- [12] Kevin P. Murphy. A survey of POMDP solution techniques. 2000.
- [13] Joseph A. O’Sullivan. Alternating minimization algorithms: from Blahut-Arimoto to expectation-maximization. In *Codes, Curves, and Signals*, pages 173–192. Springer, 1998.
- [14] Joelle Pineau, Geoffrey Gordon, and Sebastian Thrun. Point-based value iteration: an anytime algorithm for POMDPs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1025–1032, 2003.
- [15] Pascal Poupart and Craig Boutilier. Bounded finite state controllers. In *Advances in neural information processing systems*, page None, 2003.
- [16] Kenneth Rose. Deterministic annealing for clustering, compression, classification, regression, and related optimization problems. *Proceedings of the IEEE*, 86(11):2210–2239, 1998.
- [17] Nicholas Roy, Geoffrey J. Gordon, and Sebastian Thrun. Finding approximate POMDP solutions through belief compression. *Journal of Artificial Intelligence Research (JAIR)*, 23:1–40, 2005.
- [18] Jonathan Rubín, Ohad Shamir, and Naftali Tishby. Trading value and information in MDPs. In *Decision Making with Imperfect Decision Makers*, pages 57–74, 2011.
- [19] Guy Shani, Joelle Pineau, and Robert Kaplow. A survey of point-based POMDP solvers. *Autonomous Agents and Multi-Agent Systems (AAMAS)*, 27(1):1–51, 2013.
- [20] Satinder P. Singh, Tommi Jaakkola, and Michael I. Jordan. Learning without state-estimation in partially observable Markovian decision processes. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 284–292, 1994.
- [21] Steven H. Strogatz. *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering*. Perseus publishing, 2001.
- [22] Richard S. Sutton and Andrew G. Barto. *Introduction to reinforcement learning*. MIT Press, 1998.
- [23] Sekhar Tatikonda, Anant Sahai, and Sanjoy Mitter. Control of LQG systems under communication constraints. In *Proceedings of the IEEE Conference on Decision and Control*, volume 1, pages 1165–1170, 1998.
- [24] Naftali Tishby and Daniel Polani. Information theory of decisions and actions. In *Perception-Action Cycle*, pages 601–636. Springer, 2011.
- [25] Emanuel Todorov. Linearly-solvable Markov decision problems. In *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, pages 1369–1376, 2006.