
Convergence of the Bayesian Control Rule

Abstract

Recently, new approaches to adaptive control have sought to reformulate the problem as a minimization of a relative entropy criterion to obtain tractable solutions. In particular, it has been shown that minimizing the expected deviation from the causal input-output dependencies of the true plant leads to a new promising stochastic control rule called the Bayesian control rule. This work proves the convergence of the Bayesian control rule under two sufficient assumptions: boundedness, which is an ergodicity condition; and consistency, which is an instantiation of the sure-thing principle.

Keywords: Adaptive behavior, Intervention calculus, Bayesian control, Kullback-Leibler-divergence

1 Introduction

When the behavior of a plant under any control signal is fully known, then the designer can choose a controller that produces the desired dynamics. Instances of this problem include hitting a target with a cannon under known weather conditions, solving a maze having its map and controlling a robotic arm in a manufacturing plant. However, when the behavior of the plant is unknown, then the designer faces the problem of adaptive control. For example, shooting the cannon lacking the appropriate measurement equipment, finding the way out of an unknown maze and designing an autonomous robot for Martian exploration. Adaptive control turns out to be far more difficult than its non-adaptive counterpart. Even when the plant dynamics is known to belong to a particular class for which optimal controllers are available, constructing the corresponding optimal adaptive controller is in general intractable even for simple toy problems (Duff, 2002).

Thus, virtually all of the effort of the research community is centered around the development of tractable approximations.

Recently, new formulations of the adaptive control problem that are based on the minimization of a relative entropy criterion have attracted the interest of the control and reinforcement learning community. For example, it has been shown that a large class of optimal control problems can be solved very efficiently if the problem statement is reformulated as the minimization of the deviation of the dynamics of a controlled system from the uncontrolled system (Todorov, 2006, 2009; Kappen et al., 2009). A similar approach minimizes the deviation of the causal input/output-relationship of a Bayesian mixture of controllers from the true controller, obtaining an explicit solution called the *Bayesian control rule* (Ortega & Braun, 2010). This control rule is particularly interesting because it leads to stochastic controllers that infer the optimal controller on-line by combining the plant-specific controllers, implicitly using the uncertainty of the dynamics to trade-off exploration versus exploitation.

Although the Bayesian control rule constitutes a promising approach to adaptive control, there are currently no proofs that guarantee its convergence to the desired policy. The aim of this paper is to develop a set of sufficient conditions of convergence and then to provide a proof. The analysis is limited to the simple case of controllers having a finite amount of modes of operation. Special care has been taken to illustrate the motivation behind the concepts.

2 Preliminaries

The exposition is restricted to the case of discrete time with discrete stochastic observations and control signals. Let \mathcal{O} and \mathcal{A} be two finite sets of symbols, where the former is the set of inputs (observations) and the second the set of outputs (actions). Actions and observations at time t are denoted as $a_t \in \mathcal{A}$ and $o_t \in \mathcal{O}$ re-

spectively, and the shorthand $a_{\leq t} := a_1, a_2, \dots, a_t$ and the like are used to simplify the notation of strings. Symbols are underlined to glue them together as in $\underline{aO}_{\leq t} = a_1, o_1, a_2, o_2, \dots, a_t, o_t$. It is assumed that the interaction between the controller and the plant proceeds in cycles $t = 1, 2, \dots$ where in cycle t the controller issues action a_t and the plant responds with an observation o_t .

A *controller* is defined as a probability distribution P over the input/output (I/O) stream, and it is fully characterized by the conditional probabilities

$$P(a_t | \underline{aO}_{< t}) \quad \text{and} \quad P(o_t | \underline{aO}_{< t} a_t)$$

representing the probabilities of emitting action a_t and collecting observation o_t given the respective I/O history. Similarly, a *plant* is defined as a probability distribution Q characterized by the conditional probabilities

$$Q(o_t | \underline{aO}_{< t} a_t)$$

representing the probabilities of emitting observation o_t given the I/O history.

If the plant is known, i.e. if the conditional probabilities $Q(o_t | \underline{aO}_{< t} a_t)$ are known, then the designer can build a suitable controller by equating the observation streams as $P(o_t | \underline{aO}_{< t} a_t) = Q(o_t | \underline{aO}_{< t} a_t)$ and by defining action probabilities $P(a_t | \underline{aO}_{< t})$ such that the resulting distribution P maximizes a desired utility criterion. In this case P is said to be *tailored to* Q . In many situations the conditional probabilities $P(a_t | \underline{aO}_{< t})$ will be deterministic, but there are cases (e.g. in repeated games) where the designer might prefer stochastic policies instead.

If the plant is unknown then one faces an adaptive control problem. Assume we know that the plant Q_m is going to be drawn randomly from a set $\mathcal{Q} := \{Q_m\}_{m \in \mathcal{M}}$ of possible plants indexed by \mathcal{M} . Assume further we have available a set of controllers $\mathcal{P} := \{P_m\}_{m \in \mathcal{M}}$, where each P_m is tailored to Q_m . How can we now construct a controller P such that its behavior is as close as possible to the tailored controller P_m under any realization of $Q_m \in \mathcal{Q}$?

3 Bayesian Control Rule

A naïve approach would be to minimize the relative entropy of the controller P with respect to the true controller P_m , averaged over all possible values of m . However, this is syntactically incorrect. The important observation made in Ortega & Braun (2010) is that we do not want to minimize the deviation of P from P_m , but the deviation of the causal I/O dependencies in P from the causal I/O dependencies in P_m .

Intuitively speaking, one does not want to predict actions and observations, but to predict the observations (effect) given actions (causes). More specifically, they propose to minimize a set of (causal) divergences C defined by

$$C := \limsup_{t \rightarrow \infty} \sum_m P(m) \sum_{\tau=1}^t C_\tau \quad (1)$$

where

$$C_\tau := \sum_{o_{< \tau}} P_m(\hat{a}o_{< \tau}) C_\tau(\hat{a}o_{< \tau})$$

$$C_\tau(h) := \sum_{a_\tau} \sum_{o_\tau} P_m(\underline{aO}_\tau | h) \log \frac{P_m(\underline{aO}_\tau | h)}{P(\underline{aO}_\tau | h)},$$

and where $P(m)$ is the prior probability of $m \in \mathcal{M}$, \hat{a}_τ denotes an intervened (not observed) action at time τ , and $\hat{a}_1, \hat{a}_2, \hat{a}_3, \dots$ is an arbitrary sequence of intervened actions that gives rise to a particular instantiation of C .

In Ortega & Braun (2010), it is shown that the controller P that minimizes C in Equation (1) for any sequence of intervened actions is given by the conditional probabilities

$$P(a_t | \hat{a}o_{< \tau}) := \sum_m P_m(a_t | \underline{aO}_{< \tau}) P(m | \hat{a}o_{< \tau})$$

$$P(o_t | \hat{a}o_{< \tau}) := \sum_m P_m(o_t | \underline{aO}_{< \tau} a_\tau) P(m | \hat{a}o_{< \tau}) \quad (2)$$

where

$$P(m | \hat{a}o_{\leq t}) := \frac{P_m(o_t | \underline{aO}_{< t} a_t) P(m | \hat{a}o_{< t})}{\sum_{m'} P_{m'}(o_t | \underline{aO}_{< t} a_t) P(m' | \hat{a}o_{< t})}. \quad (3)$$

Equations (2) and (3) constitute the *Bayesian control rule*. This result is obtained by using properties of interventions using causal calculus (Pearl, 2000). It is worth to point out that the resulting controller is fully defined in terms of its constituent controllers in \mathcal{P} . It is customary to use the notation

$$P(a_t | m, \underline{aO}_{< t}) := P_m(a_t | \underline{aO}_{< t})$$

$$P(o_t | m, \underline{aO}_{< t} a_t) := P_m(o_t | \underline{aO}_{< t} a_t),$$

that is, treating the different controllers as “hypotheses” of a Bayesian model. In the context of the Bayesian control rule, these “I/O hypotheses” are called *operation modes*. Note that the resulting control law is in general stochastic.

4 Policy Diagrams

A policy diagram is a useful informal tool to analyze the effect of control policies on plants. Figure 1, illustrates an example. One can imagine a plant as a




Figure 1: A policy diagram.

collection of states connected by transitions labeled by I/O symbols. For instance, Figure 1 highlights a state s where taking action $a \in \mathcal{A}$ and collecting observation $o \in \mathcal{O}$ leads to state s' . In a policy diagram, one abstracts away from the underlying details of the plant's dynamics, representing sets of states and transitions as enclosed areas similar to a Venn diagram. Choosing a particular policy in a plant amounts to partially controlling the transitions taken in the state space, thereby choosing a subset of the plant's dynamics. Accordingly, a policy is represented by a subset in state space (enclosed by a directed curve) as illustrated in Figure 1.

Policy diagrams are especially useful to analyze the effect of policies on different hypotheses about the plant's dynamics. A controller that is endowed with a set of operation modes \mathcal{M} can be seen as having *hypotheses* about the plant's underlying dynamics, given by the observation models $P(o_t|m, \underline{ao}_{<t}, a_t)$, and associated *policies*, given by the action models $P(a_t|m, \underline{ao}_{<t})$, for all $m \in \mathcal{M}$. For the sake of simplifying the interpretation of policy diagrams, we will assume¹ the existence of a state space \mathcal{S} and a function $T : (\mathcal{A} \times \mathcal{O}) \rightarrow \mathcal{S}$ mapping I/O histories into states. With this assumption, policies and hypotheses can be seen as conditional probabilities

$$P(a_t|m, s) := P(a_t|m, \underline{ao}_{<t})$$

$$\text{and } P(o_t|m, s, a_t) := P(o_t|m, \underline{ao}_{<t}, a_t)$$

respectively, defining transition probabilities

$$P(s'|m, s) = \sum_{S'} P(\underline{ao}_t|m, s)$$

for a Markov chain in the state space, where $s = T(\underline{ao}_{<t})$ and S' contains the transitions \underline{ao}_t such that $T(\underline{ao}_{<t}, \underline{ao}_t) = s'$.

5 Divergence Processes

One of the obvious questions to ask oneself with respect to the Bayesian control rule is whether it converges to the right control law or not. That is, whether

¹Note however that no such assumptions are made to obtain the results of this paper.

$P(a_t|\hat{\underline{ao}}_t) \rightarrow P(a_t|m^*, \underline{ao}_{<t})$ as $t \rightarrow \infty$ when m^* is the true operation mode, i.e. the operation mode such that $P(a_t|m^*, \underline{ao}_{<t}) = Q(a_t|\underline{ao}_{<t})$. As will be obvious from the discussion in the rest of this paper, this is in general not true.

As it is easily seen from Equation 2, showing convergence amounts to show that the posterior distribution $P(m|\hat{\underline{ao}}_{<t})$ concentrates its probability mass on a subset of operation modes \mathcal{M}^* having essentially the same output stream as m^* ,

$$\sum_{m \in \mathcal{M}} P(a_t|m, \underline{ao}_{<t}) P(m|\hat{\underline{ao}}_{<t})$$

$$\approx \sum_{m \in \mathcal{M}^*} P(a_t|m^*, \underline{ao}_{<t}) P(m|\hat{\underline{ao}}_{<t})$$

$$\approx P(a_t|m^*, \underline{ao}_{<t}).$$




Figure 2: Realization of the divergence processes 1 to 4 associated to a controller with operation modes m_1 to m_4 . The divergence processes 1 and 2 diverge, whereas 3 and 4 stay below the dotted bound. Hence, the posterior probabilities of m_1 and m_2 vanish.

Hence, understanding the asymptotic behavior of the posterior probabilities

$$P(m|\hat{\underline{ao}}_{<t})$$

is the main goal of this paper. In particular, one wants to understand under what conditions these quantities converge to zero. The posterior can be rewritten as

$$P(m|\hat{\underline{ao}}_{<t}) = \frac{P(\hat{\underline{ao}}_{<t}|m)P(m)}{\sum_{m' \in \mathcal{M}} P(\hat{\underline{ao}}_{<t}|m')P(m')}$$

$$= \frac{P(m) \prod_{\tau=1}^t P(o_\tau|m, \underline{ao}_{<\tau}, a_\tau)}{\sum_{m' \in \mathcal{M}} P(m') \prod_{\tau=1}^t P(o_\tau|m', \underline{ao}_{<\tau}, a_\tau)}.$$

If all the summands but the one with index m^* are dropped from the denominator, one obtains the bound

$$P(m|\hat{\underline{ao}}_{<t}) \leq \ln \frac{P(m)}{P(m^*)} \prod_{\tau=1}^t \frac{P(o_\tau|\underline{ao}_{<\tau}, a_\tau|m)}{P(o_\tau|\underline{ao}_{<\tau}, a_\tau|m^*)},$$

which is valid for all $m \in \mathcal{M}$. From this inequality, it is seen that it is convenient to analyze the behavior

of the stochastic process

$$d_t(m^*||m) := \sum_{\tau=1}^t \ln \frac{P(o_\tau|m^*, \underline{aO}_{<\tau} a_\tau)}{P(o_\tau|m, \underline{aO}_{<\tau} a_\tau)}$$

which is the *divergence process* of m from the reference m^* . Indeed, if $d_t(m^*||m) \rightarrow \infty$ as $t \rightarrow \infty$, then

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{P(m)}{P(m^*)} \prod_{\tau=1}^t \frac{P(o_\tau|\underline{aO}_{<\tau} a_\tau|m)}{P(o_\tau|\underline{aO}_{<\tau} a_\tau|m^*)} \\ = \lim_{t \rightarrow \infty} \frac{P(m)}{P(m^*)} \cdot e^{-d_t(m^*||m)} = 0, \end{aligned}$$

and thus clearly $P(m|\hat{aO}_{\leq t}) \rightarrow 0$. Figure 2 illustrates simultaneous realizations of the divergence processes of a controller. Intuitively speaking, these processes provide lower bounds on accumulators of surprise value measured in information units.




Figure 3: The application of different policies lead to different statistical properties of the same divergence process.

A divergence process is a random walk, i.e. whose value at time t depends on the whole history up to time $t - 1$. What makes them cumbersome to characterize is the fact that their statistical properties depend on the particular policy that is applied; hence, a given divergence process can have different growth rates depending on the policy (Figure 3). Indeed, the behavior of a divergence process might depend critically on the distribution over actions that is used. For example, it can happen that a divergence process stays stable under one policy, but diverges under another. In the context of the Bayesian control rule this problem is further aggravated, because in each time step, the policy to apply is determined stochastically. More specifically, if m^* is the true operation mode, then $d_t(m^*||m)$ is a random variable that depends on the realization $\underline{aO}_{\leq t}$ which is drawn from

$$\prod_{\tau=1}^t P(a_\tau|m_\tau, \underline{aO}_{\leq \tau}) P(o_\tau|m^*, \underline{aO}_{\leq \tau} a_\tau),$$

where the m_1, m_2, \dots, m_t are drawn themselves from $P(m_1), P(m_2|\hat{aO}_1), \dots, P(m_t|\hat{aO}_{<t})$.

To deal with the heterogeneous nature of divergence processes, one can introduce a temporal decomposition that demultiplexes the original process into many sub-processes belonging to unique policies. Let $\mathcal{N}_t := \{1, 2, \dots, t\}$ be the set of time steps up to time t . Let $\mathcal{T} \subset \mathcal{N}_t$, and let $m, m' \in \mathcal{M}$. Define a *sub-divergence* of $d_t(m||m)$ as a random variable

$$g(m'; \mathcal{T}) := \sum_{\tau \in \mathcal{T}} \ln \frac{P(o_\tau|m^*, \underline{aO}_{<\tau} a_\tau)}{P(o_\tau|m, \underline{aO}_{<\tau} a_\tau)}$$

drawn from

$$\begin{aligned} P_{m'}^m(\{\underline{aO}_\tau\}_{\tau \in \mathcal{T}}|\{\underline{aO}_\tau\}_{\tau \in \mathcal{T}^c}) \\ := \left(\prod_{\tau \in \mathcal{T}} P(a_\tau|m, \underline{aO}_{<\tau}) \right) \left(\prod_{\tau \in \mathcal{T}^c} P(o_\tau|m', \underline{aO}_{<\tau} a_\tau) \right), \end{aligned}$$

where $\mathcal{T}^c := \mathcal{N}_t \setminus \mathcal{T}$ and where $\{\underline{aO}_\tau\}_{\tau \in \mathcal{T}^c}$ are given conditions that are kept constant. In this definition, m' plays the role of the policy that is used to sample the actions in the time steps \mathcal{T} . Clearly, any realization of the divergence process $d_t(m^*||m)$ can be decomposed into a sum of sub-divergences, i.e.

$$d_t(m^*||m) = \sum_{m'} g(m'; \mathcal{T}_{m'}), \quad (4)$$

where $\{\mathcal{T}_m\}_{m \in \mathcal{M}}$ forms a partition of \mathcal{N}_t . Figure 4 shows an example decomposition.




Figure 4: Decomposition of a divergence process (1) into sub-divergences (2 & 3).

The averages of sub-divergences will play an important rôle in the analysis. Define the average over all realizations of $g(m'; \mathcal{T})$ as

$$\begin{aligned} G(m', \mathcal{T}) \\ := \sum_{(\underline{aO}_\tau)_{\tau \in \mathcal{T}}} P_{m'}^m(\{\underline{aO}_\tau\}_{\tau \in \mathcal{T}}|\{\underline{aO}_\tau\}_{\tau \in \mathcal{T}^c}) g(m'; \mathcal{T}). \end{aligned}$$

Notice that for any $\tau \in \mathcal{N}_t$,

$$\begin{aligned} G(m'; \{\tau\}) \\ = \sum_{\underline{aO}_\tau} P(a_\tau|m', \underline{aO}_{<\tau}) P(o_\tau|m^*, \underline{aO}_{<\tau} a_\tau) \\ \cdot \ln \frac{P(o_\tau|m^*, \underline{aO}_{<\tau} a_\tau)}{P(o_\tau|m, \underline{aO}_{<\tau} a_\tau)} \geq 0, \end{aligned}$$

because of Gibbs' inequality. In particular,

$$G(m^*; \{\tau\}) = 0.$$

Clearly, this holds as well for any $\mathcal{T} \subset \mathcal{N}_t$:

$$\begin{aligned} \forall m' \quad G(m'; \mathcal{T}) &\geq 0, \\ G(m^*; \mathcal{T}) &= 0. \end{aligned} \quad (5)$$

6 Boundedness

In general, a divergence process is very complex: virtually all the classes of distributions that are of interest in control go well beyond i.i.d. and stationary processes. This increased complexity can jeopardize the analytic tractability of the divergence process, i.e. such that no predictions about its asymptotic behavior can be made anymore. More specifically, if the growth rates of the divergence processes vary too much from realization to realization, then the posterior distribution over operation modes can vary qualitatively between realizations. Hence, one needs to impose a stability requirement akin to ergodicity to limit the class of possible divergence-processes to a class that is analytically tractable. In the light of this insight, the following property is introduced.

A divergence process $d_t(m^*||m)$ is said to be *bounded* in \mathcal{M} iff for any $\delta > 0$, there is a $C \geq 0$, such that for all $m' \in \mathcal{M}$, all t and all $\mathcal{T} \subset \mathcal{N}_t$

$$\left| g(m'; \mathcal{T}) - G(m'; \mathcal{T}) \right| \leq C$$

with probability $\geq 1 - \delta$.




Figure 5: If a divergence process is bounded, then the realizations (curves 2 & 3) of a sub-divergence stay within a band around the mean (curve 1).

Figure 5 illustrates this property. Boundedness is the key property that is going to be used to construct the results of this paper. The first important result is that the posterior probability of the true operation mode is bounded from below.

Theorem 1. *Let the set of operation modes of a controller be such that for all $m \in \mathcal{M}$ the divergence process $d_t(m^*||m)$ is bounded. Then, for any $\delta > 0$, there*

is a $\lambda > 0$, such that for all $t \in \mathbb{N}$,

$$P(m^*|\hat{a}_{\leq t}) \geq \frac{\lambda}{|\mathcal{M}|}$$

with probability $\geq 1 - \delta$.

Proof. As has been pointed out in (4), a particular realization of the divergence process $d_t(m^*||m)$ can be decomposed as

$$d_t(m^*||m) = \sum_{m'} g_m(m'; \mathcal{T}_{m'}),$$

where the $g_m(m'; \mathcal{T}_{m'})$ are sub-divergences of $d_t(m^*||m)$ and the $\mathcal{T}_{m'}$ form a partition of \mathcal{N}_t . However, since $d_t(m^*||m)$ is bounded in \mathcal{M} , one has for all $\delta' > 0$, there is a $C(m) \geq 0$, such that for all $m' \in \mathcal{M}$, all $t \in \mathcal{N}_t$ and all $\mathcal{T} \subset \mathcal{N}_t$, the inequality

$$\left| g_m(m'; \mathcal{T}_{m'}) - G_m(m'; \mathcal{T}_{m'}) \right| \leq C(m)$$

holds with probability $\geq 1 - \delta'$. However, due to (5),

$$G_m(m'; \mathcal{T}_{m'}) \geq 0$$

for all $m' \in \mathcal{M}$. Thus,

$$g_m(m'; \mathcal{T}_{m'}) \geq -C(m).$$

If all the previous inequalities hold simultaneously then the divergence process can be bounded as well. That is, the inequality

$$d_t(m^*||m) \geq -MC(m) \quad (6)$$

holds with probability $\geq (1 - \delta')^M$ where $M := |\mathcal{M}|$. Choose

$$\beta(m) := \max\{0, \ln \frac{P(m)}{P(m^*)}\}.$$

Since $0 \geq \ln \frac{P(m)}{P(m^*)} - \beta(m)$, it can be added to the right hand side of (6). Using the definition of $d_t(m^*||m)$, taking the exponential and rearranging the terms one obtains

$$\begin{aligned} P(m^*) \prod_{\tau=1}^t P(o_\tau|m^*, \underline{a}_{\leq \tau}, a_\tau) \\ \geq e^{-\alpha(m)} P(m) \prod_{\tau=1}^t P(o_\tau|m^*, \underline{a}_{\leq \tau}, a_\tau) \end{aligned}$$

where $\alpha(m) := MC(m) + \beta(m) \geq 0$. Identifying the posterior probabilities of m^* and m by dividing both sides by the normalizing constant yields the inequality

$$P(m^*|\hat{a}_{\leq t}) \geq e^{-\alpha(m)} P(m|\hat{a}_{\leq t}).$$

This inequality holds simultaneously for all $m \in \mathcal{M}$ with probability $\geq (1 - \delta')^{M^2}$ and in particular for $\lambda := \min_m \{e^{-\alpha(m)}\}$, that is,

$$P(m^*|\hat{\underline{a}}_{\leq t}) \geq \lambda P(m|\hat{\underline{a}}_{\leq t}).$$

But since this is valid for any $m \in \mathcal{M}$, and because $\max_m \{P(m|\hat{\underline{a}}_{\leq t})\} \geq \frac{1}{M}$, one gets

$$P(m^*|\hat{\underline{a}}_{\leq t}) \geq \frac{\lambda}{M},$$

with probability $\geq 1 - \delta$ for arbitrary $\delta > 0$ related to δ' through the equation $\delta' := 1 - \sqrt[M^2]{1 - \delta}$. \square

7 Core

If one wants to identify the operation modes whose posterior probabilities vanish, then it is not enough to characterize them as those whose hypothesis does not match the true hypothesis. Figure 6 illustrates this problem. Here, three hypotheses along with their associated policies are shown. H_1 and H_2 share the prediction made for region A but differ in region B . Hypothesis H_3 differs everywhere from the others. Assume H_1 is true. As long as we apply policy P_2 , hypothesis H_3 will make wrong predictions and thus its divergence process will diverge as expected. However, no evidence against H_2 will be accumulated. It is only when we apply policy P_1 for long enough time that the controller will eventually enter region B and hence accumulate counter-evidence for H_2 .




Figure 6: If hypothesis H_1 is true and agrees with H_2 on region A , then policy P_2 cannot disambiguate the three hypotheses.

But what does “long enough” mean? If P_1 is executed only for a short period, then the controller risks not visiting the disambiguating region. But unfortunately, neither the right policy nor the right length of the period to run it are known beforehand. Hence, the controller needs a clever time-allocating strategy to test all policies for all finite time intervals. This motivates following definition.

The *core* of an operation mode m^* , denoted as $[m^*]$, is the subset of \mathcal{M} containing operation modes behaving like m^* under its policy. More formally, an operation

mode $m \notin [m^*]$ (i.e. is *not* in the core) iff for any $C \geq 0$, $\delta, \xi > 0$, there is a $t_0 \in \mathbb{N}$, such that for all $t \geq t_0$,

$$G(m^*; \mathcal{T}) \geq C$$

with probability $\geq 1 - \delta$, where $G(m^*; \mathcal{T})$ is a sub-divergence of $d_t(m^*||m)$, and $\Pr\{\tau \in \mathcal{T}\} \geq \xi$ for all $\tau \in \mathcal{N}_t$.

In other words, if the controller was to apply m^* 's policy in each time step with probability at least ξ , and under this strategy the expected sub-divergence $G(m^*; \mathcal{T})$ of $d_t(m^*||m)$ grows unboundedly, then m is not in the core of m^* . Note that demanding a strictly positive probability of execution in each time step guarantees that controller will run m^* for all possible finite time-intervals. As the following theorem shows, the posterior probabilities of the operation modes that are not in the core vanish almost surely.

Theorem 2. *Let the set of operation modes of a controller be such that for all $m \in \mathcal{M}$ the divergence process $d_t(m^*||m)$ is bounded. Then, if $m \notin [m^*]$, then $P(m|\hat{\underline{a}}_{\leq t}) \rightarrow 0$ as $t \rightarrow \infty$ almost surely.*

Proof. The divergence process $d_t(m^*||m)$ can be decomposed into a sum of sub-divergences (see Equation 4)

$$d_t(m^*||m) = \sum_{m'} g(m'; \mathcal{T}_{m'}). \quad (7)$$

Furthermore, for every $m' \in \mathcal{M}$, one has that for all $\delta > 0$, there is a $C \geq 0$, such that for all $t \in \mathbb{N}$ and for all $\mathcal{T} \subset \mathcal{N}_t$

$$\left| g(m'; \mathcal{T}) - G(m'; \mathcal{T}) \right| \leq C(m)$$

with probability $\geq 1 - \delta'$. Applying this bound to the summands in (7) yields the lower bound

$$\sum_{m'} g(m'; \mathcal{T}_{m'}) \geq \sum_{m'} (G(m'; \mathcal{T}_{m'}) - C(m))$$

which holds with probability $\geq (1 - \delta')^M$, where $M := |\mathcal{M}|$. Due to Inequality 5, one has that for all $m' \neq m^*$, $G(m'; \mathcal{T}_{m'}) \geq 0$. Hence,

$$\sum_{m'} (G(m'; \mathcal{T}_{m'}) - C(m)) \geq G(m^*; \mathcal{T}_{m^*}) - MC$$

where $C := \max_m \{C(m)\}$. The members of the set \mathcal{T}_{m^*} are determined stochastically; more specifically, the i^{th} member is included into \mathcal{T}_{m^*} with probability $P(m^*|\hat{\underline{a}}_{\leq i})$. But since $m \notin [m^*]$, one has that $G(m^*; \mathcal{T}_{m^*}) \rightarrow \infty$ as $t \rightarrow \infty$ with probability $\geq 1 - \delta'$ for arbitrarily chosen $\delta' > 0$. This implies that

$$\lim_{t \rightarrow \infty} d_t(m^*||m) \geq \lim_{t \rightarrow \infty} G(m^*; \mathcal{T}_{m^*}) - MC \nearrow \infty$$

with probability $\geq 1 - \delta$, where $\delta > 0$ is arbitrary and related to δ' as $\delta = 1 - (1 - \delta')^{M+1}$. Using this result in the upper bound for posterior probabilities yields the final result

$$0 \leq \lim_{t \rightarrow \infty} P(m|\hat{\underline{a}}_{\leq t}) \leq \lim_{t \rightarrow \infty} \frac{P(m)}{P(m^*)} e^{-d_t(m^*||m)} = 0.$$

□

8 Consistency

Even if an operation mode m is in the core of m^* , i.e. given that m is essentially indistinguishable from m^* under m^* 's control, it can still happen that m^* and m have different policies. Figure 7 shows an example of this. The hypotheses H_1 and H_2 share region A but differ in region B . In addition, both operation modes have their policies P_1 and P_2 respectively confined to region A . Note that both operation modes are in the core of each other. However, their policies are different. This means that it is unclear whether multiplexing the policies in time will ever disambiguate the two hypotheses. This is undesirable, as it could impede the convergence to the right control law.




Figure 7: An example of inconsistent policies. Both operation modes are in the core of each other, but have different policies.

Thus, it is clear that one needs to impose further restrictions on the mapping of hypotheses into policies. With respect to Figure 7, one can make the following observations:

1. Both operation modes have policies that select subsets of region A . Therefore, the dynamics in A are preferred over the dynamics in B .
2. Knowing that the dynamics in A are preferred over the dynamics in B allows to drop region B from the analysis when choosing a policy.
3. Since both hypotheses agree in region A , they have to choose the same policy in order to be consistent in their selection criterion.

This motivates the following definition. An operation mode m is said to be *consistent* with m^* iff $m \in [m^*]$

implies that for all $\varepsilon < 0$, there is a t_0 , such that for all $t \geq t_0$ and all $\underline{a} \underline{O}_{< t} a_t$,

$$\left| P(a_t|m^*, \underline{a} \underline{O}_{< t}) - P(a_t|m, \underline{a} \underline{O}_{< t}) \right| < \varepsilon.$$

In other words, if m is in the core of m^* , then m 's policy has to converge to m^* 's policy. Intuitively, this property parallels the well-known *sure-thing principle* of expected utility theory (Savage, 1954). The following theorem shows that consistency is a sufficient condition for convergence to the right control law.

Theorem 3. *Let the set of operation modes of a controller be such that: for all $m \in \mathcal{M}$ the divergence process $d_t(m^*||m)$ is bounded; and for all $m, m' \in \mathcal{M}$, m is consistent with m' . Then,*

$$P(a_t|\hat{\underline{a}}_{\leq t}) \rightarrow P(a_t|m^*, \underline{a} \underline{O}_{\leq t})$$

almost surely as $t \rightarrow \infty$.

Proof. We will use the abbreviations $p_m(t) := P(a_t|m, \hat{\underline{a}}_{\leq t})$ and $w_m(t) := P(m|\underline{a} \underline{O}_{\leq t})$. Decompose $P(a_t|\hat{\underline{a}}_{\leq t})$ as

$$P(a_t|\hat{\underline{a}}_{\leq t}) = \sum_{m \notin [m^*]} p_m(t)w_m(t) + \sum_{m \in [m^*]} p_m(t)w_m(t). \quad (8)$$

The first sum on the right-hand side is lower-bounded by zero and upper-bounded by

$$\sum_{m \notin [m^*]} p_m(t)w_m(t) \leq \sum_{m \notin [m^*]} w_m(t)$$

because $p_m(t) \leq 1$. Due to Theorem 2, $w_m(t) \rightarrow 0$ as $t \rightarrow \infty$ almost surely. Given $\varepsilon' > 0$ and $\delta' > 0$, let $t_0(m)$ be the time such that for all $t \geq t_0(m)$, $w_m(t) < \varepsilon'$. Choosing $t_0 := \max_m \{t_0(m)\}$, the previous inequality holds for all m and $t \geq t_0$ simultaneously with probability $\geq (1 - \delta')^M$. Hence,

$$\sum_{m \notin [m^*]} p_m(t)w_m(t) \leq \sum_{m \notin [m^*]} w_m(t) < M\varepsilon'. \quad (9)$$

To bound the second sum in (8) one proceeds as follows. For every member $m \in [m^*]$, one has that $p_m(t) \rightarrow p_{m^*}(t)$ as $t \rightarrow \infty$. Hence, following a similar construction as above, one can choose t'_0 such that for all $t \geq t'_0$ and $m \in [m^*]$, the inequalities

$$\left| p_m(t) - p_{m^*}(t) \right| < \varepsilon'$$

hold simultaneously for the precision $\varepsilon' > 0$. Applying

this to the first sum yields the bounds

$$\begin{aligned} & \sum_{m \in [m^*]} (p_{m^*}(t) - \varepsilon') w_m(t) \\ & \leq \sum_{m \in [m^*]} p_m(t) w_m(t) \\ & \leq \sum_{m \in [m^*]} (p_{m^*}(t) + \varepsilon') w_m(t). \end{aligned}$$

Here $(p_{m^*}(t) \pm \varepsilon')$ are multiplicative constants that can be placed in front of the sum. Note that

$$1 \geq \sum_{m \in [m^*]} w_m(t) = 1 - \sum_{m \notin [m^*]} w_m(t) > 1 - \varepsilon.$$

Diligently using of the above inequalities allows simplifying the lower and upper bounds respectively:

$$\begin{aligned} (p_{m^*}(t) - \varepsilon') \sum_{m \in [m^*]} w_m(t) &> p_{m^*}(t)(1 - \varepsilon') - \varepsilon' \\ &\geq p_{m^*}(t) - 2\varepsilon', \\ (p_{m^*}(t) + \varepsilon') \sum_{m \in [m^*]} w_m(t) &\leq p_{m^*}(t) + \varepsilon' \\ &< p_{m^*}(t) + 2\varepsilon'. \end{aligned} \tag{10}$$

Combining the inequalities (9) and (10) in (8) yields the final result:

$$\left| P(a_t | \hat{a}_{O \leq t}) - p_{m^*}(t) \right| < 3\varepsilon' = \varepsilon,$$

which holds with probability $\geq 1 - \delta$ for arbitrary $\delta > 0$ related to δ' as $\delta' = 1 - \sqrt[N]{1 - \delta}$ and arbitrary precision ε . \square

9 Summary and Conclusions

The Bayesian control rule constitutes a promising approach to adaptive control based on the minimization of the relative entropy of the causal I/O distribution of a mixture controller from the true controller. In this work, a proof of convergence of the Bayesian control rule to the true controller is provided.

Analyzing the asymptotic behavior of a controller-plant dynamics could be perceived as a difficult problem that involves the consideration of domain-specific assumptions. Here it is shown that this is not the case: the asymptotic analysis can be recast as the study of concurrent *divergence processes* that determine the evolution of the posterior probabilities over operation modes, thus abstracting away from the details of the classes of I/O distributions. In particular, if the set of operation modes is finite, then two extra assumptions are sufficient to prove convergence. The

first one, *boundedness*, imposes the stability of divergence processes under the partial influence of the policies contained within the set of operation modes. This condition can be regarded as an ergodicity assumption. The second one, *consistency*, requires that if a hypothesis makes the same predictions as another hypothesis within its most relevant subset of dynamics, then both hypotheses share the same policy. This relevance is formalized as the *core* of an operation mode.

The concepts and proof strategies developed in this work are appealing due to their intuitive interpretation and formal simplicity. Most importantly, they strengthen the intuition about potential pitfalls that arise in the context of controller design. The approach presented in this work can also be considered as a guide for possible extensions to infinite sets of operation modes. For example, one can think of partitioning a continuous space of operation modes into “essentially different” regions where representative operation modes subsume their neighborhoods (Grünwald, 2007).

Finally, convergence proofs play a crucial rôle in the mathematical justification of any new theory of control. Hopefully, this proof will contribute to establish relative entropy control theories as solid alternative formulations to the problem of adaptive control.

References

- Duff, M.O. *Optimal learning: computational procedures for bayes-adaptive markov decision processes*. PhD thesis, 2002. Director-Andrew Barto.
- Grünwald, P. *The Minimum Description Length Principle*. The MIT Press, 2007.
- Kappen, B., Gomez, V., and Opper, M. Optimal control as a graphical model inference problem. *arXiv:0901.0633*, 2009.
- Ortega, P.A. and Braun, D.A. A bayesian rule for adaptive control based on causal interventions. In *Proceedings of the third conference on general artificial intelligence*, 2010.
- Pearl, J. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2000.
- Savage, L.J. *The Foundations of Statistics*. John Wiley and Sons, New York, 1954. ISBN 0-486-62349-1.
- Todorov, E. Linearly solvable markov decision problems. In *Advances in Neural Information Processing Systems*, volume 19, pp. 1369–1376, 2006.
- Todorov, E. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences U.S.A.*, 106:11478–11483, 2009.