

A General Framework for Interacting Bayes-Optimally with Self-Interested Agents Using Arbitrary Parametric Model and Model Prior

Trong Nghia Hoang and Kian Hsiang Low

Department of Computer Science, National University of Singapore
 Republic of Singapore
 {nghiaht, lowkh}@comp.nus.edu.sg

Abstract

Recent advances in Bayesian reinforcement learning (BRL) have shown that Bayes-optimality is theoretically achievable by modeling the environment’s latent dynamics using Flat-Dirichlet-Multinomial (FDM) prior. In self-interested multi-agent environments, the transition dynamics are mainly controlled by the other agent’s stochastic behavior for which FDM’s independence and modeling assumptions do not hold. As a result, FDM does not allow the other agent’s behavior to be generalized across different states nor specified using prior domain knowledge. To overcome these practical limitations of FDM, we propose a generalization of BRL to integrate the general class of parametric models and model priors, thus allowing practitioners’ domain knowledge to be exploited to produce a fine-grained and compact representation of the other agent’s behavior. Empirical evaluation shows that our approach outperforms existing multi-agent reinforcement learning algorithms.

1 Introduction

In reinforcement learning (RL), an agent faces a dilemma between acting optimally with respect to the current, possibly incomplete knowledge of the environment (i.e., exploitation) vs. acting sub-optimally to gain more information about it (i.e., exploration). Model-based Bayesian reinforcement learning (BRL) circumvents such a dilemma by considering the notion of Bayes-optimality [Duff, 2003]: A Bayes-optimal policy selects actions that maximize the agent’s expected utility with respect to all possible sequences of future beliefs (starting from the initial belief) over candidate models of the environment. Unfortunately, due to the large belief space, the Bayes-optimal policy can only be approximately derived under a simple choice of models and model priors. For example, the Flat-Dirichlet-Multinomial (FDM) prior [Poupart *et al.*, 2006] assumes the next-state distributions for each action-state pair to be modeled as independent multinomial distributions with separate Dirichlet priors. Despite its common use to analyze and benchmark algorithms, FDM can perform poorly in practice as it often fails to exploit the struc-

tured information of a problem [Asmuth and Littman, 2011; Araya-Lopez *et al.*, 2012].

To elaborate, a critical limitation of FDM lies in its independence assumption driven by computational convenience rather than scientific insight. We can identify practical examples in the context of self-interested multi-agent RL (MARL) where the uncertainty in the transition model is mainly caused by the stochasticity in the other agent’s behavior (in different states) for which the independence assumption does not hold (e.g., motion behavior of pedestrians [Natarajan *et al.*, 2012a; 2012b]). Consider, for example, an application of BRL in the problem of placing static sensors to monitor an environmental phenomenon: It involves actively selecting sensor locations (i.e., states) for measurement such that the sum of predictive variances at the unobserved locations is minimized. Here, the phenomenon is the “other agent” and the measurements are its actions. An important characterization of the phenomenon is that of the spatial correlation of measurements between neighboring locations/states [Low *et al.*, 2007; 2008; 2009; 2011; 2012; Chen *et al.*, 2012; Cao *et al.*, 2013], which makes FDM-based BRL extremely ill-suited for this problem due to its independence assumption.

Secondly, despite its computational convenience, FDM does not permit generalization across states [Asmuth and Littman, 2011], thus severely limiting its applicability in practical problems with a large state space where past observations only come from a very limited set of states. Interestingly, in such problems, it is often possible to obtain prior domain knowledge providing a more “parsimonious” structure of the other agent’s behavior, which can potentially resolve the issue of generalization. For example, consider using BRL to derive a Bayes-optimal policy for an autonomous car to navigate successfully among human-driven vehicles [Hoang and Low, 2012; 2013b] whose behaviors in different situations (i.e., states) are governed by a small, consistent set of latent parameters, as demonstrated in the empirical study of Gipps [1981]. By estimating/learning these parameters, it is then possible to generalize their behaviors across different states. This, however, contradicts the independence assumption of FDM; in practice, ignoring this results in an inferior performance, as shown in Section 4. Note that, by using parameter tying [Poupart *et al.*, 2006], FDM can be modified to make the other agent’s behavior identical in different states. But, this simple generalization is too restrictive for real-world

problems like the examples above where the other agent's behavior in different states is not necessarily identical but related via a common set of latent "non-Dirichlet" parameters.

Consequently, there is still a huge gap in putting BRL into practice for interacting with self-interested agents of unknown behaviors. To the best of our knowledge, this is first investigated by Chalkiadakis and Boutilier [2003] who offer a myopic solution in the belief space instead of solving for a Bayes-optimal policy that is non-myopic. Their proposed BPVI method essentially selects actions that jointly maximize a heuristic aggregation of myopic value of perfect information [Dearden *et al.*, 1998] and an average estimation of expected utility obtained from solving the exact MDPs with respect to samples drawn from the posterior belief of the other agent's behavior. Moreover, BPVI is restricted to work only with Dirichlet priors and multinomial likelihoods (i.e., FDM), which are subject to the above disadvantages in modeling the other agent's behavior. Also, BPVI is demonstrated empirically in the simplest of settings with only a few states.

Furthermore, in light of the above examples, the other agent's behavior often needs to be modeled differently depending on the specific application. Grounding in the context of the BRL framework, either the domain expert struggles to best fit his prior knowledge to the supported set of models and model priors or the agent developer has to re-design the framework to incorporate a new modeling scheme. Arguably, there is no free lunch when it comes to modeling the other agent's behavior across various applications. To cope with this difficulty, the BRL framework should ideally allow a domain expert to freely incorporate his choice of design in modeling the other agent's behavior.

Motivated by the above practical considerations, this paper presents a novel generalization of BRL, which we call *Interactive BRL* (I-BRL) (Section 3), to integrate any parametric model and model prior of the other agent's behavior (Section 2) specified by domain experts, consequently yielding two advantages: The other agent's behavior can be represented (a) in a fine-grained manner based on the practitioners' prior domain knowledge, and (b) compactly to be generalized across different states, thus overcoming the limitations of FDM. We show how the non-myopic Bayes-optimal policy can be derived analytically by solving I-BRL exactly (Section 3.1) and propose an approximation algorithm to compute it efficiently in polynomial time (Section 3.2). Empirically, we evaluate the performance of I-BRL against that of BPVI [Chalkiadakis and Boutilier, 2003] using an interesting traffic problem modeled after a real-world situation (Section 4).

2 Modeling the Other Agent

In our proposed Bayesian modeling paradigm, the opponent's¹ behavior is modeled as a set of probabilities $p_{sh}^v(\lambda) \triangleq \Pr(v|s, h, \lambda)$ for selecting action v in state s conditioned on the history $h \triangleq \{s_i, u_i, v_i\}_{i=1}^d$ of d latest interactions where u_i is the action taken by our agent in the i -th step. These distributions are parameterized by λ , which abstracts the ac-

tual parametric form of the opponent's behavior; this abstraction provides practitioners the flexibility in choosing the most suitable degree of parameterization. For example, λ can simply be a set of multinomial distributions $\lambda \triangleq \{\theta_{sh}^v\}$ such that $p_{sh}^v(\lambda) \triangleq \theta_{sh}^v$ if no prior domain knowledge is available. Otherwise, the domain knowledge can be exploited to produce a fine-grained representation of λ ; at the same time, λ can be made compact to generalize the opponent's behavior across different states (e.g., Section 4).

The opponent's behavior can be learned by monitoring the belief $b(\lambda) \triangleq \Pr(\lambda)$ over all possible λ . In particular, the belief (or probability density) $b(\lambda)$ is updated at each step based on the history $h \circ \langle s, u, v \rangle$ of $d + 1$ latest interactions (with $\langle s, u, v \rangle$ being the most recent one) using Bayes' theorem:

$$b_{sh}^v(\lambda) \propto p_{sh}^v(\lambda) b(\lambda). \quad (1)$$

Let $\bar{s} = (s, h)$ denote an information state that consists of the current state and the history of d latest interactions. When the opponent's behavior is stationary (i.e., $d = 0$), it follows that $\bar{s} = s$. For ease of notations, the main results of our work (in subsequent sections) are presented only for the case where $d = 0$ (i.e., $\bar{s} = s$); extension to the general case just requires replacing s with $\bar{s}$. In this case, (1) can be re-written as

$$b_s^v(\lambda) \propto p_s^v(\lambda) b(\lambda). \quad (2)$$

The key difference between our Bayesian modeling paradigm and FDM [Poupart *et al.*, 2006] is that we do not require $b(\lambda)$ and $p_s^v(\lambda)$ to be, respectively, Dirichlet prior and multinomial likelihood where Dirichlet is a conjugate prior for multinomial. In practice, such a conjugate prior is desirable because the posterior b_s^v belongs to the same Dirichlet family as the prior b , thus making the belief update tractable and the Bayes-optimal policy efficient to be derived. Despite its computational convenience, this conjugate prior restricts the practitioners from exploiting their domain knowledge to design more informed priors (e.g., see Section 4). Furthermore, this turns out to be an overkill just to make the belief update tractable. In particular, we show in Theorem 1 below that, without assuming any specific parametric form of the initial prior, the posterior belief can still be tractably represented even though they are not necessarily conjugate distributions. This is indeed sufficient to guarantee and derive a tractable representation of the Bayes-optimal policy using a finite set of parameters, as shall be seen later in Section 3.1.

Theorem 1 *If the initial prior b can be represented exactly using a finite set of parameters, then the posterior b' conditioned on a sequence of observations $\{(s_i, v_i)\}_{i=1}^{n'}$ can also be represented exactly in parametric form.*

Proof Sketch. From (2), we can prove by induction on n' that

$$b'(\lambda) \propto \Phi(\lambda) b(\lambda) \quad (3)$$

$$\Phi(\lambda) \triangleq \prod_{s \in S} \prod_{v \in V} p_s^v(\lambda)^{\psi_s^v} \quad (4)$$

where $\psi_s^v \triangleq \sum_{i=1}^{n'} \delta_{sv}(s_i, v_i)$ and δ_{sv} is the Kronecker delta function that returns 1 if $s = s_i$ and $v = v_i$, and 0 otherwise².

²Intuitively, $\Phi(\lambda)$ can be interpreted as the likelihood of observing each pair (s, v) for ψ_s^v times while interacting with an opponent whose behavior is parameterized by λ .

¹For convenience, we will use the terms the "other agent" and "opponent" interchangeably from now on.

From (3), it is clear that b' can be represented by a set of parameters $\{\psi_s^v\}_{s,v}$ and the finite representation of b . Thus, belief update is performed simply by incrementing the hyper-parameter ψ_s^v according to each observation (s, v) . $\square$

3 Interactive Bayesian RL (I-BRL)

In this section, we first extend the proof techniques used in [Poupart *et al.*, 2006] to theoretically derive the agent's Bayes-optimal policy against the general class of parametric models and model priors of the opponent's behavior (Section 2). In particular, we show that the derived Bayes-optimal policy can also be represented exactly using a finite number of parameters. Based on our derivation, a naive algorithm can be devised to compute the exact parametric form of the Bayes-optimal policy (Section 3.1). Finally, we present a practical algorithm to efficiently approximate this Bayes-optimal policy in polynomial time (with respect to the size of the environment model) (Section 3.2).

Formally, an agent is assumed to be interacting with its opponent in a stochastic environment modeled as a tuple $(S, U, V, \{r_s\}, \{p_s^{uv}\}, \{p_s^v(\lambda)\}, \phi)$ where S is a finite set of states, U and V are sets of actions available to the agent and its opponent, respectively. In each stage, the immediate payoff $r_s(u, v)$ to our agent depends on the joint action $(u, v) \in U \times V$ and the current state $s \in S$. The environment then transitions to a new state s' with probability $p_s^{uv}(s') \triangleq \Pr(s'|s, u, v)$ and the future payoff (in state s') is discounted by a constant factor $0 < \phi < 1$, and so on. Finally, as described in Section 2, the opponent's latent behavior $\{p_s^v(\lambda)\}$ can be selected from the general class of parametric models and model priors, which subsumes FDM (i.e., independent multinomials with separate Dirichlet priors).

Now, let us recall that the key idea underlying the notion of Bayes-optimality [Duff, 2003] is to maintain a belief $b(\lambda)$ that represents the uncertainty surrounding the opponent's behavior λ in each stage of interaction. Thus, the action selected by the learner in each stage affects both its expected immediate payoff $\mathbb{E}_\lambda[\sum_v p_s^v(\lambda) r_s(u, v) | b]$ and the posterior belief state $b_s^v(\lambda)$, the latter of which influences its future payoff and builds in the information gathering option (i.e., exploration). As such, the Bayes-optimal policy can be obtained by maximizing the expected discounted sum of rewards $V_s(b)$:

$$V_s(b) \triangleq \max_u \sum_v \langle p_s^v, b \rangle \left(r_s(u, v) + \phi \sum_{s'} p_s^{uv}(s') V_{s'}(b_s^v) \right) \quad (5)$$

where $\langle a, b \rangle \triangleq \int_\lambda a(\lambda) b(\lambda) d\lambda$. The optimal policy for the learner is then defined as a function π^* that maps the belief b to an action u maximizing its expected utility, which can be derived by solving (5). To derive our solution, we first re-state two well-known results concerning the augmented belief-state MDP in single-agent RL [Poupart *et al.*, 2006], which also hold straight-forwardly for our general class of parametric models and model priors.

Theorem 2 *The optimal value function V^k for k steps-to-go converges to the optimal value function V for infinite horizon as $k \rightarrow \infty$: $\|V - V^{k+1}\|_\infty \leq \phi \|V - V^k\|_\infty$.*

Theorem 3 *The optimal value function $V_s^k(b)$ for k steps-to-go can be represented by a finite set Γ_s^k of α -functions:*

$$V_s^k(b) = \max_{\alpha_s \in \Gamma_s^k} \langle \alpha_s, b \rangle. \quad (6)$$

Simply put, these results imply that the optimal value V_s in (5) can be approximated arbitrarily closely by a finite set Γ_s^k of piecewise linear α -functions α_s , as shown in (6). Each α -function α_s is associated with an action u_{α_s} yielding an expected utility of $\alpha_s(\lambda)$ if the true behavior of the opponent is λ and consequently an overall expected reward $\langle \alpha_s, b \rangle$ by assuming that, starting from (s, b) , the learner selects action u_{α_s} and continues optimally thereafter. In particular, Γ_s^k and u_{α_s} can be derived based on a constructive proof of Theorem 3. However, due to limited space, we only state the constructive process below. Interested readers are referred to [Hoang and Low, 2013a] for a detailed proof. Specifically, given $\{\Gamma_s^k\}_s$ such that (6) holds for k , it follows (see [Hoang and Low, 2013a]) that

$$V_s^{k+1}(b) = \max_{u,t} \langle \alpha_s^{ut}, b \rangle \quad (7)$$

where $t \triangleq (t_{s'v})_{s' \in S, v \in V}$ with $t_{s'v} \in \{1, \dots, |\Gamma_{s'}^k|\}$, and

$$\alpha_s^{ut}(\lambda) \triangleq \sum_v p_s^v(\lambda) \left(r_s(u, v) + \phi \sum_{s'} \alpha_{s'}^{t_{s'v}}(\lambda) p_s^{uv}(s') \right) \quad (8)$$

such that $\alpha_{s'}^{t_{s'v}}$ denotes the $t_{s'v}$ -th α -function in $\Gamma_{s'}^k$. Setting $\Gamma_s^{k+1} = \{\alpha_s^{ut}\}_{u,t}$ and $u_{\alpha_s^{ut}} = u$, it follows that (6) also holds for $k+1$. As a result, the optimal policy $\pi^*(b)$ can be derived directly from these α -functions by $\pi^*(b) \triangleq u_{\alpha_s^*}$ where $\alpha_s^* = \arg \max_{\alpha_s^{ut} \in \Gamma_s^{k+1}} \langle \alpha_s^{ut}, b \rangle$. Thus, constructing Γ_s^{k+1} from the previously constructed sets $\{\Gamma_{s'}^k\}_s$ essentially boils down to an exhaustive enumeration of all possible pairs (u, t) and the corresponding application of (8) to compute α_s^{ut} . Though (8) specifies a bottom-up procedure constructing Γ_s^{k+1} from the previously constructed sets $\{\Gamma_{s'}^k\}_{s'}$ of α -functions, it implicitly requires a convenient parameterization for the α -functions that is closed under the application of (8). To complete this analytical derivation, we present a final result to demonstrate that each α -function is indeed of such parametric form. Note that Theorem 4 below generalizes a similar result proven in [Poupart *et al.*, 2006], the latter of which shows that, under FDM, each α -function can be represented by a linear combination of multivariate monomials. A practical algorithm building on our generalized result in Theorem 4 is presented in Section 3.2.

Theorem 4 *Let Φ denote a family of all functions $\Phi(\lambda)$ (4). Then, the optimal value $V_{s'}^k$ can be represented by a finite set $\Gamma_{s'}^k$ of α -functions $\alpha_{s'}^j$ for $j = 1, \dots, |\Gamma_{s'}^k|$:*

$$\alpha_{s'}^j(\lambda) = \sum_{i=1}^m c_i \Phi_i(\lambda) \quad (9)$$

where $\Phi_i \in \Phi$. So, each α -function $\alpha_{s'}^j$ can be compactly represented by a finite set of parameters $\{c_i\}_{i=1}^m$.³

³To ease readability, we abuse the notations $\{c_i, \Phi_i\}_{i=1}^m$ slightly: Each $\alpha_{s'}^j(\lambda)$ should be specified by a different set $\{c_i, \Phi_i\}_{i=1}^m$.

Proof Sketch. We will prove (9) by induction on k^4 . Supposing (9) holds for k . Setting $j = t_{s'v}$ in (9) results in

$$\alpha_{s'v}^{t_{s'v}}(\lambda) = \sum_{i=1}^m c_i \Phi_i(\lambda), \quad (10)$$

which is then plugged into (8) to yield

$$\alpha_s^{ut}(\lambda) = \sum_{v \in V} c_v \Psi_v(\lambda) + \sum_{s' \in S} \sum_{v \in V} \left(\sum_{i=1}^m c_{s'i}^v \Psi_{s'i}^v(\lambda) \right) \quad (11)$$

where $\Psi_v(\lambda) = p_s^v(\lambda)$, $\Psi_{s'i}^v(\lambda) = p_s^v(\lambda) \Phi_i(\lambda)$, and the coefficients $c_v = r_s(u, v)$ and $c_{s'i}^v = \phi p_s^{uv}(s') c_i$. It is easy to see that $\Psi_v \in \Phi$ and $\Psi_{s'i}^v \in \Phi$. So, (9) clearly holds for $k+1$. We have shown above that, under the general class of parametric models and model priors (Section 2), each α -function can be represented by a linear combination of arbitrary parametric functions in Φ , which subsume multivariate monomials used in [Poupart *et al.*, 2006]. $\square$

3.1 An Exact Algorithm

Intuitively, Theorems 3 and 4 provide a simple and constructive method for computing the set of α -functions and hence, the optimal policy. In step $k+1$, the sets Γ_s^{k+1} for all $s \in S$ are constructed using (10) and (11) from $\Gamma_{s'}^k$ for all $s' \in S$, the latter of which are computed previously in step k . When $k=0$ (i.e., base case), see the proof of Theorem 4 above (i.e., footnote 4). A sketch of this algorithm is shown below:

BACKUP($s, k+1$)

1. $\Gamma_{s,u}^* \leftarrow \left\{ g(\lambda) \triangleq \sum_v c_v \Psi_v(\lambda) \right\}$
2. $\Gamma_{s,u}^{v,s'} \leftarrow \left\{ g_j(\lambda) \triangleq \sum_{i=1}^m c_{s'i}^v \Psi_{s'i}^v(\lambda) \right\}_{j=1, \dots, |\Gamma_{s'}^k|}$
3. $\Gamma_{s,u} \leftarrow \Gamma_{s,u}^* \oplus \left(\bigoplus_{v,s'} \Gamma_{s,u}^{v,s'} \right)^5$
4. $\Gamma_s^{k+1} \leftarrow \bigcup_{u \in U} \Gamma_{s,u}$

In the above algorithm, steps 1 and 2 compute the first and second summation terms on the right-hand side of (11), respectively. Then, steps 3 and 4 construct $\Gamma_s^{k+1} = \{\alpha_s^{ut}\}_{u,t}$ using (11) over all t and u , respectively. Thus, by iteratively computing $\Gamma_s^{k+1} = \text{BACKUP}(s, k+1)$ for a sufficiently large value of k , Γ_s^{k+1} can be used to approximate V_s arbitrarily closely, as shown in Theorem 2. However, this naive algorithm is computationally impractical due to the following issues: (a) **α -function explosion** – the number of α -functions grows doubly exponentially in the planning horizon length, as derived from (7) and (8): $|\Gamma_s^{k+1}| = \mathcal{O}([\prod_{s'} |\Gamma_{s'}^k|]^{|V|} |U|)$, and (b) **parameter explosion** – the average number of parameters used to represent an α -function grows by a factor of $\mathcal{O}(|S||V|)$, as manifested in (11). The practicality of our approach therefore depends crucially on how these issues are resolved, as described next.

⁴When $k=0$, (9) can be verified by letting $c_i = 0$.

⁵ $A \oplus B = \{a + b | a \in A, b \in B\}$.

3.2 A Practical Approximation Algorithm

In this section, we introduce practical modifications of the **BACKUP** algorithm by addressing the above-mentioned issues. We first address the issue of **α -function explosion** by generalizing discrete POMDP's PBVI solver [Pineau *et al.*, 2003] to be used for our augmented belief-state MDP: Only the α -functions that yield optimal values for a sampled set of reachable beliefs $B_s = \{b_s^1, b_s^2, \dots, b_s^{|B_s|}\}$ are computed (see the modifications in steps 3 and 4 of the **PB-BACKUP** algorithm). The resulting algorithm is shown below:

PB-BACKUP($B_s = \{b_s^1, b_s^2, \dots, b_s^{|B_s|}\}, s, k+1$)

1. $\Gamma_{s,u}^* \leftarrow \left\{ g(\lambda) \triangleq \sum_v c_v \Psi_v(\lambda) \right\}$
2. $\Gamma_{s,u}^{v,s'} \leftarrow \left\{ g_j(\lambda) \triangleq \sum_{i=1}^m c_{s'i}^v \Psi_{s'i}^v(\lambda) \right\}_{j=1, \dots, |\Gamma_{s'}^k|}$
3. $\Gamma_{s,u}^i \leftarrow \left\{ g + \sum_{s',v} \arg \max_{g_j \in \Gamma_{s,u}^{v,s'}} \langle g_j, b_s^i \rangle \right\}_{g \in \Gamma_{s,u}^*}$
4. $\Gamma_s^{k+1} \leftarrow \left\{ g_i \triangleq \arg \max_{g \in \Gamma_{s,u}^i} \langle g, b_s^i \rangle \right\}_{i=1, \dots, |B_s|}$

Secondly, to address the issue of **parameter explosion**, each α -function is projected onto a fixed number of basis functions to keep the number of parameters from growing exponentially. This projection is done after each **PB-BACKUP** operation, hence always keeping the number of parameters fixed (i.e., one parameter per basis function). In particular, since each α -function is in fact a linear combination of functions in Φ (Theorem 4), it is natural to choose these basis functions from Φ^6 . Besides, it is easy to see from (3) that each sampled belief b_s^i can also be written as

$$b_s^i(\lambda) = \eta \Phi_s^i(\lambda) b(\lambda) \quad (12)$$

where b is the initial prior belief, $\eta = 1/\langle \Phi_s^i, b \rangle$, and $\Phi_s^i \in \Phi$. For convenience, these $\{\Phi_s^i\}_{i=1, \dots, |B_s|}$ are selected as basis functions. Specifically, after each **PB-BACKUP** operation, each $\alpha_s \in \Gamma_s^k$ is projected onto the function space defined by $\{\Phi_s^i\}_{i=1, \dots, |B_s|}$. This projection is then cast as an optimization problem that minimizes the squared difference $J(\alpha_s)$ between the α -function and its projection with respect to the sampled beliefs in B_s :

$$J(\alpha_s) \triangleq \frac{1}{2} \sum_{j=1}^{|B_s|} \left(\langle \alpha_s, b_s^j \rangle - \sum_{i=1}^{|B_s|} c_i \langle \Phi_s^i, b_s^j \rangle \right)^2. \quad (13)$$

This can be done analytically by letting $\frac{\partial J(\alpha_s)}{\partial c_i} = 0$ and solving for c_i , which is equivalent to solving a linear system $Ax = d$ where $x_i = c_i$, $A_{ji} = \sum_{k=1}^{|B_s|} \langle \Phi_s^i, b_s^k \rangle \langle \Phi_s^j, b_s^k \rangle$ and $d_j = \sum_{k=1}^{|B_s|} \langle \Phi_s^j, b_s^k \rangle \langle \alpha_s, b_s^k \rangle$. Note that this projection works directly with the values $\langle \alpha_s, b_s^j \rangle$ instead of the exact parametric form of α_s in (9). This allows for a more compact

⁶See Appendix B of Hoang and Low [2013a] for other choices.

implementation of the **PB-BACKUP** algorithm presented above: Instead of maintaining the exact parameters that represent each of the immediate functions g , only their evaluations at the sampled beliefs $B_s = \{b_s^1, b_s^2, \dots, b_s^{|B_s|}\}$ need to be maintained. In particular, the values of $\{\langle g, b_s^i \rangle\}_{i=1, \dots, |B_s|}$ can be estimated as follows:

$$\begin{aligned} \langle g, b_s^i \rangle &= \eta \int_{\lambda} g(\lambda) \Phi_s^i(\lambda) b(\lambda) d\lambda \\ &\approx \frac{\sum_{j=1}^n g(\lambda^j) \Phi_s^i(\lambda^j)}{\sum_{j=1}^n \Phi_s^i(\lambda^j)} \end{aligned} \quad (14)$$

where $\{\lambda^j\}_{j=1}^n$ are samples drawn from the initial prior b . During the online execution phase, (14) is also used to compute the expected payoff for the α -functions evaluated at the current belief $b'(\lambda) = \eta \Phi(\lambda) b(\lambda)$:

$$\langle \alpha_s, b' \rangle \approx \frac{\sum_{j=1}^n \Phi(\lambda^j) \sum_{i=1}^{|B_s|} c_i \Phi_s^i(\lambda^j)}{\sum_{j=1}^n \Phi(\lambda^j)}. \quad (15)$$

So, the real-time processing cost of evaluating each α -function's expected reward at a particular belief is $\mathcal{O}(|B_s|n)$. Since the sampling of $\{b_s^i\}$, $\{\lambda^j\}$ and the computation of $\{\sum_{i=1}^{|B_s|} c_i \Phi_s^i(\lambda^j)\}$ can be performed in advance, this $\mathcal{O}(|B_s|n)$ cost is further reduced to $\mathcal{O}(n)$, which makes the action selection incur $\mathcal{O}(|B_s|n)$ cost in total. This is significantly cheaper as compared to the total cost $\mathcal{O}(nk|S|^2|U||V|)$ of online sampling and re-estimating V_s incurred by BPVI [Chalkiadakis and Boutilier, 2003]. Also, note that since the offline computational costs in steps 1 to 4 of **PB-BACKUP**($B_s, s, k+1$) and the projection cost, which is cast as the cost of solving a system of linear equations, are always polynomial functions of the interested variables (e.g., $|S|, |U|, |V|, n, |B_s|$), the optimal policy can be approximated in polynomial time.

4 Experiments and Discussion

In this section, a realistic scenario of intersection navigation is modeled as a stochastic game; it is inspired from a near-miss accident during the 2007 DARPA Urban Challenge. Considering the traffic situation illustrated in Fig. 1 where two autonomous vehicles (marked A and B) are about to enter an intersection (I), the road segments are discretized into a uniform grid with cell size $5 \text{ m} \times 5 \text{ m}$ and the speed of each vehicle is also discretized uniformly into 5 levels ranging from 0 m/s to 4 m/s. So, in each stage, the system's state can be characterized as a tuple $\{P_A, P_B, S_A, S_B\}$ specifying the current positions (P) and velocities (S) of A and B, respectively. In addition, our vehicle (A) can either accelerate ($+1 \text{ m/s}^2$), decelerate (-1 m/s^2), or maintain its speed ($+0 \text{ m/s}^2$) in each time step while the other vehicle (B) changes its speed based on a parameterized reactive model [Gipps, 1981]:

$$\begin{aligned} v_{\text{safe}} &= S_B + \frac{\text{Distance}(P_A, P_B) - \tau S_B}{S_B/d + \tau} \\ v_{\text{des}} &= \min(4, S_B + a, v_{\text{safe}}) \\ S'_B &\sim \text{Uniform}(\max(0, v_{\text{des}} - \sigma a), v_{\text{des}}). \end{aligned}$$

In this model, the driver's acceleration $a \in [0.5 \text{ m/s}^2, 3 \text{ m/s}^2]$, deceleration $d \in [-3 \text{ m/s}^2, -0.5 \text{ m/s}^2]$, reaction time $\tau \in [0.5\text{s}, 2\text{s}]$, and imperfection $\sigma \in [0, 1]$ are the unknown parameters distributed uniformly within the corresponding ranges. This parameterization can cover a variety of drivers' typical behaviors, as shown in a preliminary study. For a further understanding of these parameters, the readers are referred to [Gipps, 1981]. Besides, in each time step, each vehicle $X \in \{A, B\}$ moves from its current cell P_X to the next cell P'_X with probability $1/t$ and remains in the same cell with probability $1 - 1/t$ where t is the expected time to move forward one cell from the current position with respect to the current speed (e.g., $t = 5/S_X$). Thus, in general, the underlying stochastic game has $6 \times 6 \times 5 \times 5 = 900$ states (i.e., each vehicle has 6 possible positions and 5 levels of speed), which is significantly larger than the settings in previous experiments. In each state, our vehicle has 3 actions, as mentioned previously, while the other vehicle has 5 actions corresponding to 5 levels of speed according to the reactive model.




Figure 1: (Left) A near-miss accident during the 2007 DARPA Urban Challenge, and (Right) the discretized environment: A and B move towards destinations D_A and D_B while avoiding collision at I. Shaded areas are not passable.

The goal for our vehicle in this domain is to learn the other vehicle's reactive model and adjust its navigation strategy accordingly such that there is no collision and the time spent to cross the intersection is minimized. To achieve this goal, we penalize our vehicle in each step by -1 and reward it with 50 when it successfully crosses the intersection. If it collides with the other vehicle (at I), we penalize it by -250 . The discount factor is set as 0.99. We evaluate the performance of I-BRL in this problem against 100 different sets of reactive parameters (for the other vehicle) generated uniformly from the above ranges. Against each set of parameters, we run 20 simulations ($h = 100$ steps each) to estimate our vehicle's average performance⁷ R . In particular, we compare our algorithm's average performance against the average performance of a fully informed vehicle (**Upper Bound**) who knows exactly the reactive parameters before each simulation, a rational vehicle (**Exploit**) who estimates the reactive parameters by taking the means of the above ranges, and a vehicle employing BPVI [Chalkiadakis and Boutilier, 2003] (**BPVI**).

The results are shown in Fig. 2a: It can be observed that our vehicle always performs significantly better than both the rational and BPVI-based vehicles. In particular, our vehicle manages to reduce the performance gap between the

⁷After our vehicle successfully crosses the intersection, the system's state is reset to the default state in Fig. 1 (Right).




Figure 2: (a) Performance comparison between our vehicle (I-BRL), the fully informed, the rational and the BPVI vehicles ($\phi = 0.99$); (b) Our approach's offline planning time.

fully informed and rational vehicles roughly by half. The difference in performance between our vehicle and the fully informed vehicle is expected as the fully informed vehicle always takes the optimal step from the beginning (since it knows the reactive parameters in advance) while our vehicle has to take cautious steps (by maintaining a slow speed) before it feels confident with the information collected during interaction. Intuitively, the performance gap is mainly caused during this initial period of “caution”. Also, since the uniform prior over the reactive parameters $\lambda = \{a, d, \tau, \sigma\}$ is not a conjugate prior for the other vehicle's behavior model $\theta_s(v) = p(v|s, \lambda)$, the BPVI-based vehicle has to directly maintain and update its belief using FDM: $\lambda = \{\theta_s\}_s$ with $\theta_s = \{\theta_s^v\}_v \sim \text{Dir}(\{n_s^v\}_v)$ (Section 2), instead of $\lambda = \{a, d, \tau, \sigma\}$. However, FDM implicitly assumes that $\{\theta_s\}_s$ are statistically independent, which is not true in this case since all θ_s are actually related by $\{a, d, \tau, \sigma\}$. Unfortunately, BPVI cannot exploit this information to generalize the other vehicle's behavior across different states due to its restrictive FDM (i.e., independent multinomial likelihoods with separate Dirichlet priors), thus resulting in an inferior performance.

5 Related Works

In self-interested (or non-cooperative) MARL, there has been several groups of proponents advocating different learning goals, the following of which have garnered substantial support: (a) **Stability** – in self-play or against a certain class of learning opponents, the learners' behaviors converge to an equilibrium; (b) **optimality** – a learner's behavior necessarily converges to the best policy against a certain class of learning opponents; and (c) **security** – a learner's average payoff must exceed the maximin value of the game. For example, the works of Littman [2001], Bianchi *et al.* [2007], and Akchurina [2009] have focused on (evolutionary) game-theoretic approaches that satisfy the **stability** criterion in self-play. The works of Bowling and Veloso [2001], Suematsu and Hayashi [2002], and Tesauro [2003] have developed algorithms that address both the **optimality** and **stability** criteria: A learner essentially converges to the best response if the opponents' policies are stationary; otherwise, it converges in self-play. Notably, the work of Powers and Shoham [2005] has proposed an approach that provably converges to an ϵ -best response (i.e., **optimality**) against a class of adaptive, bounded-memory opponents while simultaneously guaran-

teeing a minimum average payoff (i.e., **security**) in single-state, repeated games.

In contrast to the above-mentioned works that focus on convergence, I-BRL directly optimizes a learner's performance during its course of interaction, which may terminate before it can successfully learn its opponent's behavior. So, our main concern is how well the learner can perform before its behavior converges. From a practical perspective, this seems to be a more appropriate goal: In reality, the agents may only interact for a limited period, which is not enough to guarantee convergence, thus undermining the **stability** and **optimality** criteria. In such a context, the existing approaches appear to be at a disadvantage: (a) Algorithms that focus on **stability** and **optimality** tend to select exploratory actions with drastic effect without considering their huge costs (i.e., poor rewards) [Chalkiadakis and Boutilier, 2003]; and (b) though the notion of **security** aims to prevent a learner from selecting such radical actions, the proposed security values (e.g., maximin value) may not always turn out to be tight lower bounds for the optimal performance [Hoang and Low, 2013a]. Interested readers are referred to [Chalkiadakis and Boutilier, 2003] and Appendix C of Hoang and Low [2013a] for a detailed discussion and additional experiments to compare performances of I-BRL and these approaches, respectively.

Note that while solving for the Bayes-optimal policy efficiently has not been addressed explicitly in general prior to this paper, we can actually avoid this problem by allowing the agent to act sub-optimally in a bounded number of steps. In particular, the works of Asmuth and Littman [2011] and Araya-Lopez *et al.* [2012] guarantee that, in the worst case, the agent will act nearly approximately Bayes-optimal in all but a polynomially bounded number of steps with high probability. It is thus necessary to point out the difference between I-BRL and these worst-case approaches: We are interested in maximizing the average-case performance with certainty rather than the worst-case performance with some “high probability” guarantee. Comparing their performances is beyond the scope of this paper.

6 Conclusion

This paper describes a novel generalization of BRL, called I-BRL, to integrate the general class of parametric models and model priors of the opponent's behavior. As a result, I-BRL relaxes the restrictive assumption of FDM that is often imposed in existing works, thus offering practitioners greater flexibility in encoding their prior domain knowledge of the opponent's behavior. Empirical evaluation shows that I-BRL outperforms a Bayesian MARL approach utilizing FDM called BPVI. I-BRL also outperforms existing MARL approaches focusing on convergence (Section 5), as shown in the additional experiments in [Hoang and Low, 2013a]. To this end, we have successfully bridged the gap in applying BRL to self-interested multi-agent settings.

Acknowledgments. This work was supported by Singapore-MIT Alliance Research & Technology Subaward Agreements No. 28 R-252-000-502-592 & No. 33 R-252-000-509-592.

References

- [Akchurina, 2009] Natalia Akchurina. Multiagent reinforcement learning: Algorithm converging to Nash equilibrium in general-sum discounted stochastic games. In *Proc. AAMAS*, pages 725–732, 2009.
- [Araya-Lopez *et al.*, 2012] M. Araya-Lopez, V. Thomas, and O. Buffet. Near-optimal BRL using optimistic local transitions. In *Proc. ICML*, 2012.
- [Asmuth and Littman, 2011] J. Asmuth and M. L. Littman. Learning is planning: Near Bayes-optimal reinforcement learning via Monte-Carlo tree search. In *Proc. UAI*, pages 19–26, 2011.
- [Bianchi *et al.*, 2007] Reinaldo A. C. Bianchi, Carlos H. C. Ribeiro, and Anna H. R. Costa. Heuristic selection of actions in multiagent reinforcement learning. In *Proc. IJCAI*, 2007.
- [Bowling and Veloso, 2001] Michael Bowling and Manuela Veloso. Rational and convergent learning in stochastic games. In *Proc. IJCAI*, 2001.
- [Cao *et al.*, 2013] N. Cao, K. H. Low, and J. M. Dolan. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In *Proc. AAMAS*, 2013.
- [Chalkiadakis and Boutilier, 2003] Georgios Chalkiadakis and Craig Boutilier. Coordination in multiagent reinforcement learning: A Bayesian approach. In *Proc. AAMAS*, pages 709–716, 2003.
- [Chen *et al.*, 2012] J. Chen, K. H. Low, C. K.-Y. Tan, A. Oran, P. Jaillet, J. M. Dolan, and G. S. Sukhatme. Decentralized data fusion and active sensing with mobile sensors for modeling and predicting spatiotemporal traffic phenomena. In *Proc. UAI*, pages 163–173, 2012.
- [Dearden *et al.*, 1998] Richard Dearden, Nir Friedman, and Stuart Russell. Bayesian Q-learning. In *Proc. AAAI*, pages 761–768, 1998.
- [Duff, 2003] Michael Duff. Design for an optimal probe. In *Proc. ICML*, 2003.
- [Gipps, 1981] P. G. Gipps. A behavioural car following model for computer simulation. *Transportation Research B*, 15(2):105–111, 1981.
- [Hoang and Low, 2012] T. N. Hoang and K. H. Low. Intention-aware planning under uncertainty for interacting with self-interested, boundedly rational agents. In *Proc. AAMAS*, pages 1233–1234, 2012.
- [Hoang and Low, 2013a] T. N. Hoang and K. H. Low. A general framework for interacting Bayes-optimally with self-interested agents using arbitrary parametric model and model prior. arXiv:1304.2024, 2013.
- [Hoang and Low, 2013b] T. N. Hoang and K. H. Low. Interactive POMDP Lite: Towards practical planning to predict and exploit intentions for interacting with self-interested agents. In *Proc. IJCAI*, 2013.
- [Littman, 2001] Michael L. Littman. Friend-or-foe Q-learning in general-sum games. In *Proc. ICML*, 2001.
- [Low *et al.*, 2007] K. H. Low, G. J. Gordon, J. M. Dolan, and P. Khosla. Adaptive sampling for multi-robot wide-area exploration. In *Proc. IEEE ICRA*, 2007.
- [Low *et al.*, 2008] K. H. Low, J. M. Dolan, and P. Khosla. Adaptive multi-robot wide-area exploration and mapping. In *Proc. AAMAS*, pages 23–30, 2008.
- [Low *et al.*, 2009] K. H. Low, J. M. Dolan, and P. Khosla. Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing. In *Proc. ICAPS*, pages 233–240, 2009.
- [Low *et al.*, 2011] K. H. Low, J. M. Dolan, and P. Khosla. Active Markov information-theoretic path planning for robotic environmental sensing. In *Proc. AAMAS*, pages 753–760, 2011.
- [Low *et al.*, 2012] K. H. Low, J. Chen, J. M. Dolan, S. Chien, and D. R. Thompson. Decentralized active robotic exploration and mapping for probabilistic field classification in environmental sensing. In *Proc. AAMAS*, pages 105–112, 2012.
- [Natarajan *et al.*, 2012a] P. Natarajan, T. N. Hoang, K. H. Low, and M. Kankanhalli. Decision-theoretic approach to maximizing observation of multiple targets in multi-camera surveillance. In *Proc. AAMAS*, pages 155–162, 2012.
- [Natarajan *et al.*, 2012b] P. Natarajan, T. N. Hoang, K. H. Low, and M. Kankanhalli. Decision-theoretic coordination and control for active multi-camera surveillance in uncertain, partially observable environments. In *Proc. ICDSC*, 2012.
- [Pineau *et al.*, 2003] J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. In *Proc. IJCAI*, pages 1025–1032, 2003.
- [Poupart *et al.*, 2006] Pascal Poupart, Nikos Vlassis, Jesse Hoey, and Kevin Regan. An analytic solution to discrete Bayesian reinforcement learning. In *Proc. ICML*, pages 697–704, 2006.
- [Powers and Shoham, 2005] Rob Powers and Yoav Shoham. Learning against opponents with bounded memory. In *Proc. IJCAI*, 2005.
- [Suematsu and Hayashi, 2002] Nobuo Suematsu and Akira Hayashi. A multiagent reinforcement learning algorithm using extended optimal response. In *Proc. AAMAS*, 2002.
- [Tesauro, 2003] Gerald Tesauro. Extending Q-learning to general adaptive multi-agent systems. In *Proc. NIPS*, 2003.