

Feature Reinforcement Learning: Part II. Structured MDPs*

Marcus Hutter

Australian National University
Canberra, ACT, 2600, Australia

WWW.HUTTERI.NET

Editor: Alexey Potapov

Abstract

The Feature Markov Decision Processes (Φ MDPs) model developed in Part I (Hutter, 2009b) is well-suited for learning agents in general environments. Nevertheless, unstructured (Φ)MDPs are limited to relatively simple environments. Structured MDPs like Dynamic Bayesian Networks (DBNs) are used for large-scale real-world problems. In this article I extend Φ MDP to Φ DBN. The primary contribution is to derive a cost criterion that allows to automatically extract the most relevant features from the environment, leading to the “best” DBN representation. I discuss all building blocks required for a complete general learning algorithm, and compare the novel Φ DBN model to the prevalent POMDP approach.

Keywords: Reinforcement learning; dynamic Bayesian network; structure learning; feature selection; global vs. local reward; explore-exploit; information & complexity; rational agents; partial observability.

1. Introduction

Agents. The agent-environment setup in which an *agent* interacts with an *environment* is a very general and prevalent framework for studying intelligent learning systems (Russell and Norvig, 2003). In cycles $t = 1, 2, 3, \dots$, the environment provides a (regular) *observation* $o_t \in \mathcal{O}$ (e.g. a camera image) to the agent; then the agent chooses an *action* $a_t \in \mathcal{A}$ (e.g. a limb movement); finally the environment provides a real-valued *reward* $r_t \in \mathbb{R}$ to the agent. The reward may be very scarce, e.g. just +1 (-1) for winning (losing) a chess game, and 0 at all other times (Hutter, 2005, Sec.6.3). Then the next cycle $t+1$ starts. The agent’s objective is to maximize his reward.

Environments. For example, *sequence prediction* is concerned with environments that do not react to the agent’s actions (e.g. a weather-forecasting “action”) (Hutter, 2003), *planning* deals with the case where the environmental function is known (Ross et al., 2008), *classification* and *regression* is for conditionally independent observations (Bishop, 2006), *Markov Decision Processes* (MDPs) assume that o_t and r_t only depend on a_{t-1} and o_{t-1} (Puterman, 1994; Sutton and Barto, 2018), POMDPs deal with *Partially Observable MDPs* (Kaelbling, Littman, and Cassandra, 1998), and *Dynamic Bayesian Networks* (DBNs) with structured MDPs (Boutilier, Dean, and Hanks, 1999).

*, Originally I had greater ambitions for this Part II, but I never found the time, so I finally decided to publish the version last edited on 9th of June 2009 as is. It extends the conference version (Hutter, 2009a), which won the AGI Kurzweil award back then, by Sections 8&9&A.

Feature MDPs (Hutter, 2009b). Concrete real-world problems can often be *modeled* as MDPs. For this purpose, a *designer* extracts relevant features from the history (e.g. position and velocity of all objects), i.e. the *history* $h_t = o_1 a_1 r_1 \dots o_{t-1} a_{t-1} r_{t-1} o_t$ is summarized by a *feature* vector $s_t := \Phi(h_t)$. The feature vectors are regarded as *states* of an MDP and are assumed to be (approximately) Markov.

Artificial General Intelligence (AGI) (Goertzel and Pennachin, 2007) is concerned with designing *agents that perform well in a very large range of environments* (Legg and Hutter, 2007), including all of the mentioned ones above and more. In this general situation, it is not a priori clear what the useful features are. Indeed, any observation in the (far) past may be relevant in the future. A solution suggested by Hutter (2009b) is to learn Φ itself.

If Φ keeps too much of the history (e.g. $\Phi(h_t) = h_t$), the resulting MDP is too large (infinite) and cannot be learned. If Φ keeps too little, the resulting state sequence is not Markov. The *Cost* criterion I develop formalizes this tradeoff and is minimized for the “best” Φ . At any time n , the best Φ is the one that minimizes the Markov code length of $s_1 \dots s_n$ and $r_1 \dots r_n$. This reminds but is actually quite different from MDL, which minimizes model+data code length (Grünwald, 2007).

Dynamic Bayesian networks. The use of “unstructured” MDPs (Hutter, 2009b), even our Φ -optimal ones, is clearly limited to relatively simple tasks. Real-world problems are structured and can often be represented by dynamic Bayesian networks (DBNs) with a reasonable number of nodes (Dean and Kanazawa, 1989). Bayesian networks in general and DBNs in particular are powerful tools for modeling and solving complex real-world problems. Advances in theory and increase in computation power constantly broaden their range of applicability (Boutilier, Dean, and Hanks, 1999; Strehl, Diuk, and Littman, 2007).

Main contribution. The primary contribution of this work is to extend the Φ *selection principle* developed by Hutter (2009b) for MDPs to the conceptually much more demanding DBN case. The major extra complications are approximating, learning and coding the rewards, the dependence of the Cost criterion on the DBN structure, learning the DBN structure, and how to store and find the optimal value function and policy.

Contents. Section 2 formalizes our Φ DBN setup, which consists of the agent model with a map Φ from observation-reward-action histories to states of a structured MDP. Section 3 explains Φ DBN on a simple example. Section 4 develops our core Φ selection principle for DBNs. Section 5 discusses search algorithms for finding (approximations of) the optimal Φ by adding and removing features. I also briefly discusses searching for the best graphical DBNs structure. Section 6 reviews state of the art algorithms for finding optimal actions in DBN, including exploration. Section 7 shows how a significant speedup can be obtained by incrementally updating various quantities of interest with time. Section 8 compares Φ DBN to POMDPs, the prevalent models for dealing with complex (partially observable) environments. Section 10 contains a necessarily brief outlook.

Although this article is self-contained, it is recommended to read Part I (Hutter, 2009b) first.

2. Feature Dynamic Bayesian Networks (Φ DBN)

In this section I recapitulate the definition of Φ MDP from Part I, and adapt it to DBNs. While formally a DBN is just a special case of an MDP, exploiting the additional structure efficiently is a challenge. For generic MDPs, typical algorithms should be polynomial and can at best be linear in the number of states $|S|$. For DBNs we want algorithms that are polynomial in the number of features m . Such DBNs have exponentially many states ($2^{O(m)}$), hence the standard MDP

algorithms are exponential, not polynomial, in m . Deriving poly-time (and poly-space!) algorithms for DBNs by exploiting the additional DBN structure is the challenge.

The gain is that we can handle exponentially large structured MDPs efficiently.

Notation. Throughout this article, $\log$ denotes the binary logarithm, and $\delta_{x,y} = \delta_{xy} = 1$ if $x = y$ and 0 else is the Kronecker symbol. I generally omit separating commas if no confusion arises, in particular in indices. For any z of suitable type (string,vector,set), I define string $\mathbf{z} = z_{1:l} = z_1 \dots z_l$, sum $z_+ = \sum_j z_j$, union $z_* = \bigcup_j z_j$, and vector $\mathbf{z}_\bullet = (z_1, \dots, z_l)$, where j ranges over the full range $\{1, \dots, l\}$ and $l = |z|$ is the length or dimension or size of z . $\hat{z}$ denotes an estimate of z . The characteristic function $\mathbb{1}_B = 1$ if $B=\text{true}$ and 0 else. $P(\cdot)$ denotes a probability over states and rewards or parts thereof. I do not distinguish between random variables Z and realizations z , and abbreviation $P(z) := P[Z = z]$ never leads to confusion. More specifically, $m \in \mathbb{N}$ denotes the number of features, $i \in \{1, \dots, m\}$ any feature, $n \in \mathbb{N}$ the current time, and $t \in \{1, \dots, n\}$ any time in history. Further, in order not to get distracted at several places I gloss over initial conditions or special cases where inessential. Also $0 * \text{undefined} = 0 * \text{infinity} = 0$.

Φ MDP definition. A Φ MDP consists of a 7 tuple $(\mathcal{O}, \mathcal{A}, \mathcal{R}, \text{Agent}, \text{Env}, \Phi, \mathcal{S}) := (\text{observation space, action space, reward space, agent, environment, feature map, state space})$. $\mathcal{A}$ is assumed to be finite and small and $\mathcal{R} \subseteq \mathbb{R}$. No assumptions are made on $\mathcal{O}$; indeed, Φ MDP has been specifically designed to cope with huge or even infinite observation spaces.

Agent and Env are a pair or triple of interlocking functions of the history $\mathcal{H} := (\mathcal{O} \times \mathcal{A} \times \mathcal{R})^* \times \mathcal{O}$:

$$\begin{aligned} \text{Env} : \mathcal{H} \times \mathcal{A} \times \mathcal{R} &\rightsquigarrow \mathcal{O}, & o_n &= \text{Env}(h_{n-1} a_{n-1} r_{n-1}), \\ \text{Agent} : \mathcal{H} &\rightsquigarrow \mathcal{A}, & a_n &= \text{Agent}(h_n), \\ \text{Env} : \mathcal{H} \times \mathcal{A} &\rightsquigarrow \mathcal{R}, & r_n &= \text{Env}(h_n a_n). \end{aligned}$$

where $\rightsquigarrow$ indicates that mappings $\rightarrow$ might be stochastic. The informal goal of AI is to design an Agent() that achieves high (expected) reward over the agent's lifetime in a large range of Env()ironments.

The feature map Φ maps histories to states

$$\Phi : \mathcal{H} \rightarrow \mathcal{S}, \quad s_t = \Phi(h_t), \quad h_t = o_1 a_1 r_1 \dots o_{t-1} a_{t-1} r_{t-1} o_t \in \mathcal{H}$$

Formally, a Φ MDP is a stochastic process Env() such that the ‘‘compressed’’ history $s_1 a_1 r_1 \dots s_n a_n r_n$ is generated by an MDP $(\mathcal{S}, \mathcal{A}, T, R)$ for some transition probability and reward function. But we are mainly interested in *finding* Φ that extract the ‘‘relevant’’ aspects of the history in the sense that the compressed history $sar_{1:n}$ can *approximately* be described as a sample from some MDP.

(Φ) Dynamic Bayesian Networks are structured (Φ)MDPs. The state space is $\mathcal{S} = \{0, 1\}^m$, and each state $s \equiv \mathbf{x} \equiv (x^1, \dots, x^m) \in \mathcal{S}$ is interpreted as a feature vector $\mathbf{x} = \Phi(h)$, where $x^i = \Phi^i(h)$ is the value of the i th binary feature. In the following I will also refer to x^i as feature i , although strictly speaking it is its value. Since non-binary features can be realized as a list of binary features, I restrict myself to the latter.

Given $\mathbf{x}_{t-1} = \mathbf{x}$, I assume that the features $(x_t^1, \dots, x_t^m) = \mathbf{x}'$ at time t are independent, and that each x^i depends only on a subset of ‘‘parent’’ features $\mathbf{u}^i \subseteq \{x^1, \dots, x^m\}$, i.e. the transition matrix has the structure

$$T_{\mathbf{x}\mathbf{x}'}^a = P(\mathbf{x}_t = \mathbf{x}' | \mathbf{x}_{t-1} = \mathbf{x}, a_{t-1} = a) = \prod_{i=1}^m P(x^i | \mathbf{u}^i, a) \quad (1)$$

This defines our **Φ DBN model**. It is just a Φ MDP with special S and T . Explaining Φ DBN on an example is easier than staying general.

3. Φ DBN Example

Consider an instantiation of the simple vacuum world (Russell and Norvig, 2003, Sec.3.6). There are two rooms, A and B , and a vacuum Robot that can observe whether the room he is in is *C*lean or *D*irty; *M*ove to the other room, *S*uck, i.e. clean the room he is in; or do *N*othing. After 3 days a room gets dirty again. Every clean room gives a reward 1, but a moving or sucking robot costs and hence reduces the reward by 1. Hence $\mathcal{O} = \{A, B\} \times \{C, D\}$, $\mathcal{A} = \{N, S, M\}$, $\mathcal{R} = \{-1, 0, 1, 2\}$, and the dynamics $\text{Env}()$ (possible histories) is clear from the above description.

Dynamics as a DBN. We can model the dynamics by a DBN as follows: The state is modeled by 3 features. Feature $R \in \{A, B\}$ stores in which room the robot is, and feature $A/B \in \{0, 1, 2, 3\}$ remembers (capped at 3) how long ago the robot has cleaned room A/B last time, hence $S = \{0, 1, 2, 3\} \times \{A, B\} \times \{0, 1, 2, 3\}$. The state/feature transition is as follows:

$$\begin{aligned} &\text{if } (x^R = A \text{ and } a = S) \text{ then } x'^A = 0 \text{ else } x'^A = \min\{x^A + 1, 3\}; \\ &\text{if } (x^R = B \text{ and } a = S) \text{ then } x'^B = 0 \text{ else } x'^B = \min\{x^B + 1, 3\}; \\ &\text{if } a = M \text{ (if } x^R = B \text{ then } x'^R = A \text{ else } x'^R = B) \text{ else } x'^R = x^R; \end{aligned}$$


A DBN can be viewed as a two-layer Bayesian network (Boutilier, Dean, and Hanks, 1999). The dependency structure of our example is depicted in the right diagram.

Each feature consists of a (left,right)-pair of nodes, and a node $i \in \{1, 2, 3 = m\} \triangleq \{A, R, B\}$ on the right is connected to all and only the parent features u^i on the left. The reward is

$$r = \mathbb{1}_{x^A < 3} + \mathbb{1}_{x^B < 3} - \mathbb{1}_{a \neq N}$$

The features map $\Phi = (\Phi^A, \Phi^R, \Phi^B)$ can also be written down explicitly. It depends on the actions and observations of the last 3 time steps.

Discussion. Note that all nodes x^i can implicitly also depend on the chosen action a . The optimal policies are repetitions of action sequence S, N, M or S, M, N . One might think that binary features $x^{A/B} \in \{C, D\}$ are sufficient, but this would result in a POMDP (Partially Observable MDP), since the cleanness of room A is not observed while the robot is in room B . That is, x' would not be a (probabilistic) function of x and a alone. This issue is discussed in more detail later. The quaternary feature $x^A \in \{0, 1, 2, 3\}$ can easily be converted into two binary features, and similarly x^B . The purely deterministic example can easily be made stochastic. For instance, *Sucking* and *Moving* may fail with a certain probability. Possible, but more complicated is to model a probabilistic transition from *C*lean to *D*irty. In the randomized versions the agent needs to use his observations.



4. Φ DBN Coding and Evaluation

I now construct a code for $s_{1:n}$ given $a_{1:n}$, and for $r_{1:n}$ given $s_{1:n}$ and $a_{1:n}$, which is optimal (minimal) if $s_{1:n}r_{1:n}$ given $a_{1:n}$ is sampled from some MDP. It constitutes our cost function for

Φ and is used to define the Φ selection principle for DBNs. Compared to the MDP case, reward coding is more complex, and there is an extra dependence on the graphical structure of the DBN.

Recall from Part I (Hutter, 2009b) that a sequence $z_{1:n}$ with counts $\mathbf{n} = (n_1, \dots, n_m)$ can within an additive constant be coded in

$$\text{CL}(\mathbf{n}) := n H(\mathbf{n}/n) + \frac{m'-1}{2} \log n \text{ if } n > 0 \text{ and } 0 \text{ else} \quad (2)$$

bits, where $n = n_+ = n_1 + \dots + n_m$ and $m' = |\{i : n_i > 0\}| \leq m$ is the number of non-empty categories, and $H(\mathbf{p}) := -\sum_{i=1}^m p_i \log p_i$ is the entropy of probability distribution $\mathbf{p}$. The code is optimal (within $+O(1)$) for all i.i.d. sources.

State/Feature Coding. Similarly to the Φ MDP case, we need to code the temporal “observed” state=feature sequence $\mathbf{x}_{1:n}$. I do this by a frequency estimate of the state/feature transition probability. (Any method of Part I leads to the same result). In the following I will drop the prime in $(\mathbf{u}^i, a, x^i)$ tuples and related situations if/since it does not lead to confusion. Let $\mathcal{T}_{\mathbf{u}^i x^i}^{ia} = \{t \leq n : \mathbf{u}_{t-1} = \mathbf{u}^i, a_{t-1} = a, x_t^i = x^i\}$ be the set of times $t-1$ at which features that influence x^i have values $\mathbf{u}^i$, and action is a , and which leads to feature i having value x^i . Let $n_{\mathbf{u}^i x^i}^{ia} = |\mathcal{T}_{\mathbf{u}^i x^i}^{ia}|$ their number ($n_{++}^{i+} = n \forall i$). I estimate each feature probability separately by $\hat{P}(x^i | \mathbf{u}^i, a) = n_{\mathbf{u}^i x^i}^{ia} / n_{\mathbf{u}^i +}^{ia}$. Using (1), this yields

$$\hat{P}(\mathbf{x}_{1:n} | a_{1:n}) = \prod_{t=1}^n \hat{T}_{\mathbf{x}_{t-1} \mathbf{x}_t}^{a_{t-1} a_t} = \prod_{t=1}^n \prod_{i=1}^m \hat{P}(x_t^i | \mathbf{u}_{t-1}^i, a_{t-1}) = \exp \left[\sum_{i, \mathbf{u}^i, a} n_{\mathbf{u}^i +}^{ia} H \left(\frac{n_{\mathbf{u}^i \bullet}^{ia}}{n_{\mathbf{u}^i +}^{ia}} \right) \right]$$

The length of the Shannon-Fano code of $\mathbf{x}_{1:n}$ is just the logarithm of this expression. We also need to code each non-zero count $n_{\mathbf{u}^i x^i}^{ia}$ to accuracy $O(1/\sqrt{n_{\mathbf{u}^i +}^{ia}})$, which each needs $\frac{1}{2} \log(n_{\mathbf{u}^i +}^{ia})$ bits. Together this gives a complete code of length

$$\text{CL}(\mathbf{x}_{1:n} | a_{1:n}) = \sum_{i, \mathbf{u}^i, a} \text{CL}(n_{\mathbf{u}^i \bullet}^{ia}) \quad (3)$$

The expression ignores the code of the graphical structure, since its length is independent of n . An alternative derivation is to use (2) and realize that for each $(i, \mathbf{u}^i, a)$, $(x_t : t \in \mathcal{T}_{\mathbf{u}^i x^i}^{ia})$ constitutes an i.i.d. sequence with counts $n_{\mathbf{u}^i x^i}^{ia}$. The rewards are more complicated.

Reward structure. Let $R_{\mathbf{x}\mathbf{x}'}^a$ be (a model of) the observed reward when action a in state $\mathbf{x}$ results in state $\mathbf{x}'$. It is natural to assume that the structure of the rewards $R_{\mathbf{x}\mathbf{x}'}^a$ is related to the transition structure $T_{\mathbf{x}\mathbf{x}'}^a$. Indeed, this is not restrictive, since one can always consider a DBN with the union of transition and reward dependencies. Usually it is assumed that the “global” reward is a *sum* of “local” rewards $R_{\mathbf{u}^i x^i}^{ia}$, one for each feature i (Koller and Parr, 1999). For simplicity of exposition I assume that the local reward R^i only depends on the feature value x^i and not on $\mathbf{u}^i$ and a . Even this is not restrictive and actually may be advantageous as discussed in Part I for MDPs. So I assume

$$R_{\mathbf{x}\mathbf{x}'}^a = \sum_{i=1}^m R_{x^i}^i =: R(\mathbf{x}')$$

For instance, in the example in the previous section, two local rewards ($R_{x^1 A}^A = \mathbb{1}_{x^1 A < 3}$ and $R_{x^2 B}^B = \mathbb{1}_{x^2 B < 3}$) depend on $\mathbf{x}'$ only, but the third reward depends on the action ($R_{a \neq N}^R = -\mathbb{1}_{a \neq N}$).

Often it is assumed that the local rewards are directly observed or known (Koller and Parr, 1999), but we neither want nor can do this here: Having to specify many local rewards is an extra burden for the environment (e.g. the teacher), which preferably should be avoided. In our case, it is not even possible to pre-specify a local reward for each feature, since the features Φ^i themselves are learned by the agent and are not statically available. They are agent-internal and not part of the Φ DBN interface. In case multiple rewards *are* available, they can be modeled as part of the regular observations o , and r only holds the overall reward. The agent must and can learn to interpret and exploit the local rewards in o by himself.

Learning the reward function. In analogy to the MDP case for R and the DBN case for T above it is tempting to estimate $R_{x^i}^i$ by $\sum_{r'} r' n_{+x^i}^{ir'}/n_{+x^i}^{i+}$ but this makes no sense. For instance if $r_t = 1 \forall t$, then $\hat{R}_{x^i}^i \equiv 1$, and $\hat{R}_{xx'}^a \equiv m$ is a gross mis-estimation of $r_t \equiv 1$. The localization of the global reward is somewhat more complicated. The goal is to choose $R_{x^1}^1, \dots, R_{x^m}^m$ such that $r_t = R(x_t) \forall t$.

Without loss we can set $R_0^i \equiv 0$, since we can subtract a constant from each local reward and absorb them into an overall constant w_0 . This allows us to write

$$R(x) = w_0 x^0 + w_1 x^1 + \dots + w_m x^m = \mathbf{w}^\top \mathbf{x}$$

where $w_i := R_1^i$ and $x^0 := 1$.

In practice, the Φ DBN model will not be perfect, and an approximate solution, e.g. a least squares fit, is the best we can achieve. The square loss can be written as

$$\begin{aligned} \text{Loss}(\mathbf{w}) &:= \sum_{t=1}^n (R(x_t) - r_t)^2 = \mathbf{w}^\top A \mathbf{w} - 2\mathbf{b}^\top \mathbf{w} + c \\ A_{ij} &:= \sum_{t=1}^n x_t^i x_t^j, \quad b_i := \sum_{t=1}^n r_t x_t^i, \quad c := \sum_{t=1}^n r_t^2 \end{aligned} \quad (4)$$

Note that A_{ij} counts the number of times feature i and j are “on” ($=1$) simultaneously, and b_i sums all rewards for which feature i is on. The loss is minimized for

$$\hat{\mathbf{w}} := \arg \min_{\mathbf{w}} \text{Loss}(\mathbf{w}) = A^{-1} \mathbf{b}, \quad \hat{R}(x) = \hat{\mathbf{w}}^\top x$$

which involves an inversion of the $(m+1) \times (m+1)$ matrix A . For singular A we take the pseudo-inverse.

Reward coding. The quadratic loss function suggests a Gaussian model for the rewards:

$$P(r_{1:n} | \hat{\mathbf{w}}, \sigma) := \exp(-\text{Loss}(\hat{\mathbf{w}})/2\sigma^2) / (2\pi\sigma^2)^{n/2}$$

Maximizing this w.r.t. the variance σ^2 yields the maximum likelihood estimate

$$-\log P(r_{1:n} | \hat{\mathbf{w}}, \hat{\sigma}) = \frac{n}{2} \log(\text{Loss}(\hat{\mathbf{w}})) - \frac{n}{2} \log \frac{ne}{2\pi}$$

where $\hat{\sigma}^2 = \text{Loss}(\hat{\mathbf{w}})/n$. Given $\hat{\mathbf{w}}$ and $\hat{\sigma}$ this can be regarded as the (Shannon-Fano) code length of $r_{1:n}$ (there are actually a few subtleties here which I gloss over). Each weight $\hat{w}_k$ and $\hat{\sigma}$ need also be coded to accuracy $O(1/\sqrt{n})$, which needs $(m+2)\frac{1}{2}\log n$ bits total. Together this gives a complete code of length

$$\text{CL}(r_{1:n} | x_{1:n} a_{1:n}) = \frac{n}{2} \log\left(\frac{2\pi}{ne} \text{Loss}(\hat{\mathbf{w}})\right) + \frac{m+2}{2} \log n \quad (5)$$

Φ DBN evaluation and selection is similar to the MDP case. Let G denote the graphical structure of the DBN, i.e. the set of parents $\text{Pa}^i \subseteq \{1, \dots, m\}$ of each feature i . (Remember u^i are the parent values). Similarly to the MDP case, the cost of (Φ, G) on h_n is defined as

$$\text{Cost}(\Phi, G|h_n) := \text{CL}(x_{1:n}|a_{1:n}) + \text{CL}(r_{1:n}|x_{1:n}, a_{1:n}), \quad (6)$$

Simplistic Φ (e.g. $\Phi \equiv 0$) lead to a short (vacuous) code for $x_{1:n}$ but obscure potential structure in the reward sequence, leading to a long code for $r_{1:n}$. Over-detailed Φ (e.g. $\Phi(h) = h$) make coding $r_{1:n}$ easy (trivial) but lead to a long code for $x_{1:n}$. Those Φ that only keep the relevant aspects of the history lead to a reasonably short code for $x_{1:n}$ and $r_{1:n}$, and the best (Φ, G) minimizes this cost.

$$(\Phi^{best}, G^{best}) := \arg \min_{\Phi, G} \{\text{Cost}(\Phi, G|h_n)\}$$

A more detailed discussion why this is a good criterion can be found in Part I. In the following section I mainly highlight the difference to the MDP case, in particular the additional dependence on and optimization over G .

5. DBN Structure Learning & Updating

This section briefly discusses minimization of (6) w.r.t. G given Φ and even briefer minimization w.r.t. Φ . For the moment regard Φ as given and fixed.

Cost and DBN structure. For general structured local rewards $R_{u^i x^i}^{ia}$, (3) and (5) both depend on G , and (6) represents a novel DBN structure learning criterion that includes the rewards.

For our simple reward model $R_{x^i}^i$, (5) is independent of G , hence only (3) needs to be considered. This is a standard MDL criterion, but I have not seen it used in DBNs before. Further, the features i are independent in the sense that we can search for the optimal parent sets $\text{Pa}^i \subseteq \{1, \dots, m\}$ for each feature i separately.

Complexity of structure search. Even in this case, finding the optimal DBN structure is generally hard. In principle we could rely on off-the-shelf heuristic search methods for finding good G , but it is probably better to use or develop some special purpose optimizer. One may even restrict the space of considered graphs G to those for which (6) can be minimized w.r.t. G efficiently, as long as this restriction can be compensated by “smarter” Φ .

A brute force exhaustive search algorithm for Pa^i is to consider all 2^m subsets of $\{1, \dots, m\}$ and select the one that minimizes $\sum_{u^i, a} \text{CL}(n_{u^i, a}^{ia})$. A reasonable and often employed assumption is to limit the number of parents to some small value p , which reduces the search space size to $O(m^p)$.

Indeed, since the Cost is exponential in the maximal number of parents of a feature, but only linear in n , a Cost minimizing Φ can usually not have more than a logarithmic number of parents, which leads to a search space that is pseudo-polynomial in m .

Heuristic structure search. We could also replace the well-founded criterion (3) by some heuristic. One such heuristic has been developed by Strehl, Diuk, and Littman (2007). The mutual information is another popular criterion for determining the dependency of two random variables, so we could add j as a parent of feature i if the mutual information of x^j and x^i is above a certain threshold. Overall this takes time $O(m^2)$ to determine G . An MDL inspired threshold for binary random variables is $\frac{1}{2n} \log n$. Since the mutual information treats parents independently, $\hat{T}$ has to be estimated accordingly, essentially as in naive Bayes classification (Lewis, 1998) with feature

selection, where x'^i represents the class label and u^i are the features selected x . The improved Tree-Augmented naive Bayes (TAN) classifier (Friedman, Geiger, and Goldszmid, 1997) could be used to model synchronous feature dependencies (i.e. within a time slice). The Chow-Liu (Chow and Liu, 1968) minimum spanning tree algorithm allows determining G in time $O(m^3)$. A tree becomes a forest if we employ a lower threshold for the mutual information.

Φ search is even harder than structure search, and remains an art. Nevertheless the reduction of the complex (ill-defined) reinforcement learning problem to an internal feature search problem with well-defined objective is a clear conceptual advance.

In principle (but not in practice) we could consider the set of *all* (computable) functions $\{\Phi : \mathcal{H} \rightarrow \{0,1\}\}$. We then compute $\text{Cost}(\Phi|h)$ for every finite subset $\Phi = \{\Phi^{i_1}, \dots, \Phi^{i_m}\}$ and take the minimum (note that the order is irrelevant).

Most practical search algorithms require the specification of some neighborhood function, here for Φ . For instance, stochastic search algorithms suggest and accept a neighbor of Φ with a probability that depends on the Cost reduction. See Part I for more details. Here I will only present some very simplistic ideas for features and neighborhoods.

Assume binary observations $\mathcal{O} = \{0,1\}$ and consider the last m observations as features, i.e. $\Phi^i(h_n) = o_{n-i+1}$ and $\Phi(h_n) = (\Phi^1(h_n), \dots, \Phi^m(h_n)) = o_{n-m+1:n}$. So the states are the same as for Φ_m MDP in Part I, but now $\mathcal{S} = \{0,1\}^m$ is structured as m binary features. In the example here, $m = 5$ lead to a perfect Φ DBN. We can add a new feature o_{n-m} ($m \rightsquigarrow m+1$) or remove the last feature ($m \rightsquigarrow m-1$), which defines a natural neighborhood structure.

Note that the context trees of McCallum (1996); Hutter (2009b) are more flexible. To achieve this flexibility here we either have to use smarter features within our framework (simply interpret $s = \Phi_{\mathcal{S}}(h)$ as a feature vector of length $m = \lceil \log|\mathcal{S}| \rceil$) or use smarter (non-tabular) estimates of $P(x^i|u^i, a)$ extending our framework (to tree dependencies).

For general purpose intelligent agents we clearly need more powerful features. Logical expressions or (non)accepting Turing machines or recursive sets can map histories or parts thereof into true/false or accept/reject or in/out, respectively, hence naturally represent binary features. Randomly generating such expressions or programs with an appropriate bias towards simple ones is a universal feature generator that eventually finds the optimal feature map. The idea is known as Universal Search (Gaglio, 2007).

6. Value & Policy Learning in Φ DBN

Given an estimate $\hat{\Phi}$ of Φ^{best} , the next step is to determine a good action for our agent. I mainly concentrate on the difficulties one faces in adapting MDP algorithms and discuss state of the art DBN algorithms. Value and policy learning in known finite state MDPs is easy provided one is satisfied with a polynomial time algorithm. Since a DBN is just a special (structured) MDP, its (Q) Value function respects the same Bellman equations

$$Q_x^{*a} = \sum_{x'} T_{xx'}^a [R_{xx'}^a + \gamma V_{x'}^*] \quad \text{and} \quad V_x^* = \max_a Q_x^{*a}$$

where $0 < \gamma < 1$ is a discount parameter, typically close to 1. See (Hutter, 2005, Sec.5.7) for proper choices. The optimal policy is also still given by

$$a_{n+1} := \arg \max_a Q_{x_{n+1}}^{*a}, \quad \text{where} \quad x_{n+1} = \Phi(h_{n+1})$$

Nevertheless, their solution is now a nightmare, since the state space is exponential in the number of features. We need algorithms that are polynomial in the number of features, i.e. logarithmic in the number of states.

Value function approximation. The first problem is that the optimal value and policy do not respect the structure of the DBN. They are usually complex functions of the (exponentially many) states, which cannot even be stored, not to mention computed (Koller and Parr, 1999). It has been suggested that the value can often be approximated well as a sum of local values similarly to the rewards. Such a value function can at least be stored.

Model-based learning. The default quality measure for the approximate value is the ρ -weighted squared difference, where ρ is the stationary distribution.

Even for a fixed policy, value iteration does *not* converge to the best approximation, but usually converges to a fixed point close to it (Bertsekas and Tsitsiklis, 1996). Value iteration requires ρ explicitly. Since ρ is also too large to store, one has to approximate ρ as well. Another problem, as pointed out by Koller and Parr (2000), is that policy iteration may not converge, since different policies have different (misleading) stationary distributions. Koller and Parr (2000) devised algorithms for general factored ρ , and Guestrin et al. (2003) for max-norm, alleviating this problem. Finally, general policies cannot be stored exactly, and another restriction or approximation is necessary.

Model-free learning. Given the difficulties above, I suggest to (re)consider a very simple class of algorithms, without suggesting that it is better. The above model-based algorithms exploit $\hat{T}$ and $\hat{R}$ directly. An alternative is to sample from $\hat{T}$ and use model-free “Temporal Difference (TD)” learning algorithms based only on this internal virtual sample (Sutton and Barto, 2018). We could use $TD(\lambda)$ or Q -value variants with linear value function approximation.

Beside their simplicity, another advantage is that neither the stationary distribution nor the policy needs to be stored or approximated. Once approximation $\hat{Q}^*$ has been obtained, it is trivial to determine the optimal (w.r.t. $\hat{Q}^*$) action via $a_{n+1} = \arg\max_a Q_{x_{n+1}}^{*a}$ for any state of interest (namely x_{n+1}) exactly.

Exploration. Optimal actions based on approximate rather than exact values can lead to very poor behavior due to lack of exploration. There are polynomially optimal algorithms (Rmax, E3, OIM) for the exploration-exploitation dilemma.

For model-based learning, extending E3 to DBNs is straightforward, but E3 needs an oracle for planning in a given DBN (Kearns and Koller, 1999). Recently, Strehl et al. (2007) accomplished the same for Rmax. They even learn the DBN structure, albeit in a very simplistic way. Algorithm OIM (Szita and Lörincz, 2008), which I described in Part I for MDPs, can also likely be generalized to DBNs, and I can imagine a model-free version.

7. Incremental Updates

As discussed two sections ago, most search algorithms are local in the sense that they produce a chain of “slightly” modified candidate solutions, here Φ ’s. This suggests a potential speedup by computing quantities of interest incrementally.

Cost. Computing $CL(x|a)$ in (3) takes at most time $O(m2^k|A|)$, where k is the maximal number of parents of a feature. If we remove feature i , we can simply remove/subtract the contributions from

i in the sum. If we add a new feature $m+1$, we only need to search for the best parent set $\mathbf{u}^{m+1}$ for this new feature, and add the corresponding code length. In practice, many transitions don't occur, i.e. $n_{\mathbf{u}^i x^i}^{ia} = 0$, so $\text{CL}(\mathbf{x}|\mathbf{a})$ can actually be computed much faster in time $O(|\{n_{\mathbf{u}^i x^i}^{ia} > 0\}|)$, and incrementally even faster.

Rewards. When adding a new feature, the current local reward estimates may not change much. If we reassign a fraction $\alpha \leq 1$ of reward to the new feature x^{m+1} , we get the following ansatz¹.

$$\begin{aligned}\hat{R}(x^1, \dots, x^{m+1}) &= (1-\alpha)\hat{R}(\mathbf{x}) + w_{m+1}x^{m+1} =: \mathbf{v}^\top \boldsymbol{\psi}(\mathbf{x}) \\ \mathbf{v} &:= (1-\alpha, w_{m+1})^\top, \quad \boldsymbol{\psi} := (\hat{R}(\mathbf{x}), x^{m+1})^\top\end{aligned}$$

Minimizing $\sum_{t=1}^n (\hat{R}(x_t^1 \dots x_t^{m+1}) - r_t)^2$ w.r.t. $\mathbf{v}$ analogous to (4) just requires a trivial 2×2 matrix inversion. The minimum $\tilde{\mathbf{v}}$ results in an initial new estimate $\tilde{\mathbf{w}} = ((1-\tilde{\alpha})\hat{w}_0, \dots, (1-\tilde{\alpha})\hat{w}_m, \tilde{w}_{m+1})^\top$, which can be improved by some first order gradient decent algorithm in time $O(m)$, compared to the exact $O(m^3)$ algorithm. When removing a feature, we simply redistribute its local reward to the other features, e.g. uniformly, followed by improvement steps that cost $O(m)$ time.

Value. All iteration algorithms described in the previous section for computing (Q) Values need an initial value for V or Q . We can take the estimate $\hat{V}$ from a previous Φ as an initial value for the new Φ . Similarly as for the rewards, we can redistribute a fraction of the values by solving relatively small systems of equations. The result is then used as an initial value for the iteration algorithms in the previous section. A further speedup can be obtained by using prioritized iteration algorithms that concentrate their time on badly estimated parameters, which are in our case the new values (Sutton and Barto, 2018).

Similarly, results from time t can be (re)used as initial estimates for the next cycle $t+1$, followed by a fast improvement step.

8. Relation to Partial observable MDPs

Partial observable MDPs. Recently, in RL, much attention has been devoted to Partially Observable MDPs (POMDPs) (Kaelbling, Littman, and Cassandra, 1998; Ross et al., 2008). Nature is assumed to be an MDP, but the states $s \in \mathcal{S}$ of nature are only partially observed via some non-injective or probabilistic function, i.e. the agent observes $o = O(s) + \text{noise}$. Typically $\mathcal{S}$ is assumed to be finite and small, mostly because already this case causes a headache.

Even for known observation and transition functions, finding the optimal action is (harder than NP) hard. The most prevalent approach is to reduce the POMDP to an MDP over so-called belief states. At every point in time, the agent maintains a belief distribution b over states, i.e. $b(s) \in [0; 1]$ is the degree to which the agent believes that he is in state s . After an observation o and action a , Bayesian updating gives the agent's new belief b' .

The difficulty starts with computing or even storing the value function, which is a convex but otherwise quite arbitrary function of the $|\mathcal{S}| - 1$ real parameters $(b(s))_{s \in \mathcal{S}}$. See Ross et al. (2008) and references therein for the state of the art.

If the transition function is unknown, it need to be learned. A simple frequency estimate is no longer possible, since the state is not observed. Nature often is a large but structured MDP, e.g. a

1. An *Ansatz* is an initial mathematical or physical model with some free parameters to be determined subsequently. [<http://en.wikipedia.org/wiki/Ansatz>]

DBN, i.e. POMDP learning algorithms need to be extended to “PODBNs”, which seems to be an untouched field. Finally, when lifting the unrealistic assumptions of a known observation function or a known number of states or ask for exploration policies we enter terra nullius.

Relation between POMDP and Φ MDP/DBN. In the following I compare the prevalent POMDP approach to our novel Φ MDP approach. Since b_{t+1} is a function of o_t and a_t and b_t , we can recurse back in time, and compute b_n from the initial belief b_0 (provided by the POMDP designer) and $o_{1:t-1}$ and $a_{1:t-1}$. Hence $b_n = f(h_n)$ is some a priori fixed function f of history h_n . Indeed, there is nothing else b_n could depend on than the history. Function f is “just” a specific complicated feature function Φ with one real-valued feature per state of nature, which may approximately be converted into (more) binary features. This is exactly what happened in the tiny example in Hutter (2009b, Sec.4). Nature was a POMDP, since the robot could not observe the cleanness of the “other” room. I transformed this into an MDP by converting the cleanness state/feature into a belief, which could in this simple example be represented by a quaternary feature.

For unknown transition and observation function, learning b is the analogue of learning Φ . Every b corresponds to some Φ . For unknown number of nature states, it is conceivable that reversely every consistent Φ corresponds to some b , e.g. by choosing $\mathcal{H}$ as nature state space. So formally, both approaches might be equivalent. But do they lead to the same “optimal” solution? This is impossible to say to date, since a general POMDP learning algorithm has yet to be suggested.

Why is the POMDP approach pursued despite all its complications? Here are a few possible answers: (1) Closed physical systems form a continuous Markov process, where the phase space is the state space. If we allow an external agent to interact with such a system, and discretize, we arrive at a POMDP. Hence, a given (narrow) problem (class) can often directly be modeled as a POMDP(s). (2) There was no real alternative: The interesting predictive state representation (Littman, Sutton, and Singh, 2001) is in an even less developed stage (Singh et al., 2003; Singh, James, and Rudary, 2004); to this date, AIXI (Hutter, 2005) is a purely theoretical alternative; and the potential alternative Φ MDP/DBN has just now been suggested. (3) Even with this alternative, while Φ -search is well-defined, *efficiently* searching for (approximations of) Φ^{best} is an open problem.

Can the two approaches profit from each other? There might be some generic properties that all interesting environments share. This property might easily be characterized in the POMDP formulation. So the POMDP metaphor may help us to restrict the space of potentially interesting features Φ (the reverse is also thinkable). Conversely, the Φ DBN cost criterion may be transferable to POMDPs, which would make (the search for) general POMDP learning (algorithms) a well-defined task.

9. Miscellaneous

Storing the history. The $\text{Cost}(\Phi|h_n)$ depends on the complete history. While in simple situations, one may restrict to features depending on the k most recent observations for some fixed k , in general, we need Φ ’s that compute statistics of (e.g. count) events for very large k , possibly $k = n$. This requires storage and access to the whole history h_n .

Toy agents for studying Φ DBN may have a (real-world equivalent) lifespan of a month with 1KB of observations per (real-world equivalent) second. This accumulates 2.5GB of data, which just fits into the main memory of a PC.

Consider now a real general purpose robot equipped with a video camera. We need about 1GB for one hour of reasonable quality mpeg movie. So over a lifespan of 100 years, the robot collects about 1PB= 10^{15} Byte of data. For 50 years now, memory increases by about a factor of 2 every 1.5 years. Assuming Moore’s law also holds in the future, 1PB fits on a home-computer’s hard-drive in 20 years and main memory in 30 years. This is also otherwise a realistic time-frame to make AGI work.

So there is no storage problem. Note that it is likely not necessary to access all history in every single time-step, hence infrequently used parts can be stored on the hard-drive. Storing only such aspects of the history that were relevant so far is risky, since any other aspect might turn out to be relevant in the future.

Dimensions of intelligence. Intelligence has many facets: creativity, problem solving, pattern recognition, classification, learning, induction, deduction, analogical reasoning, optimization, surviving in an environment, language processing, vision, information processing, knowledge, and others. It is beyond the scope of this paper to discuss to which extend Φ DBN possesses these properties. Most of the core properties are clearly incorporated in Φ DBN at least to some extent, partly by design and partly emergent. The more peripheral interface properties like vision are only indirectly addressed: If observations are sequences of images, Cost will be small for image features that are useful for reward prediction.

Possible types of environments and problems a general agent might face are: sequential (prediction), i.i.d (classification/regression), supervised, unsupervised, reinforcement learning, known/unknown environment, exploration, exploitation, planning, fully/partially observed MDPs, unstructured (MDP) versus structured (DBN), competitive (multi-agents, games) versus stochastic environment (single agent), and optimization problems. Since Φ DBN was designed to cope with general action-observation-reward environments, and all the environments above can be regarded as specializations, Φ DBN should be able to deal with them, at least to some extent. Again, a detailed analysis is beyond the scope of this article, but see Hutter (2005); Legg and Hutter (2007) for such a discussion for the AIXI model.

10. Discussion

Outlook. Φ DBN leaves much more questions open and room for modifications and improvements than Φ MDP. Here are a few.

- The cost function can be improved by integrating out the states analogous to the Φ MDP case Part I: The likelihood $P(r_{1:n}|a_{1:n}, \hat{U})$ is unchanged, except that $\hat{U} \equiv \hat{T}\hat{R}$ is now estimated locally, and the complexity penalty becomes $\frac{1}{2}(M+m+2)\log n$, where M is (essentially) the number of non-zero counts $n_{u^i x^i}^{ia}$, but an efficient algorithm has yet to be found.
- It may be necessary to impose and exploit structure on the conditional probability tables $P^a(x^i|u^i)$ themselves (Boutilier, Dean, and Hanks, 1999).
- Real-valued observations and beliefs suggest to extend the binary feature model to $[0,1]$ interval valued features rather than coding them binary. Since any continuous semantics that preserves the role of 0 and 1 is acceptable, there should be an efficient way to generalize Cost and Value estimation procedures.
- I assumed that the reward/value is linear in local rewards/values. Is this sufficient for all practical purposes? I also assumed a least squares and Gaussian model for the local rewards.

There are efficient algorithms for much more flexible models. The least we could do is to code w.r.t. the proper covariance A .

- I also barely discussed synchronous (within time-slice) dependencies.
- I guess Φ DBN will often be able to work around too restrictive DBN models, by finding features Φ that are more compatible with the DBN and reward structure.
- Extra edges in the DBN can improve the linear value function approximation. To give Φ DBN incentives to do so, the Value would have to be included in the Cost criterion.
- Implicitly I assumed that the action space $\mathcal{A}$ is small. It is possible to extend Φ DBN to large structured action spaces.
- Apart from the Φ -search, all parts of Φ DBN seem to be poly-time approximable, which is satisfactory in theory. In practice, this needs to be improved to essentially linear time in n and m .
- Developing smart Φ generation and smart stochastic search algorithms for Φ are the major open challenges.
- A more Bayesian Cost criterion would be desirable: a likelihood of h given Φ and a prior over Φ leading to a posterior of Φ given h , or so. Monte Carlo (search) algorithms like Metropolis-Hastings could sample from such a posterior. Currently probabilities ($\cong 2^{-CL}$) are assigned only to rewards and states, but not to observations and feature maps.

Summary. In this work I introduced a powerful framework (Φ DBN) for general-purpose intelligent learning agents, and presented algorithms for all required building blocks. The introduced cost criterion reduced the informal reinforcement learning problem to an internal well-defined search for “relevant” features.

References

- Bertsekas, D. P., and Tsitsiklis, J. N. 1996. *Neuro-Dynamic Programming*. Belmont, MA: Athena Scientific.
- Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*. Springer.
- Boutilier, C.; Dean, T.; and Hanks, S. 1999. Decision-Theoretic Planning: Structural Assumptions and Computational Leverage. *Journal of Artificial Intelligence Research* 11:1–94.
- Chow, C. K., and Liu, C. N. 1968. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory* IT-14(3):462–467.
- Dean, T., and Kanazawa, K. 1989. A Model for Reasoning about Persistence and Causation. *Computational Intelligence* 5(3):142–150.
- Friedman, N.; Geiger, D.; and Goldszmid, M. 1997. Bayesian Network Classifiers. *Machine Learning* 29(2):131–163.
- Gaglio, M. 2007. Universal Search. *Scholarpedia* 2(11):2575.
- Goertzel, B., and Pennachin, C., eds. 2007. *Artificial General Intelligence*. Springer.
- Grünwald, P. D. 2007. *The Minimum Description Length Principle*. Cambridge: The MIT Press.

- Guestrin, C.; Koller, D.; Parr, R.; and Venkataraman, S. 2003. Efficient Solution Algorithms for Factored MDPs. *Journal of Artificial Intelligence Research (JAIR)* 19:399–468.
- Hutter, M. 2003. Optimality of Universal Bayesian Prediction for General Loss and Alphabet. *Journal of Machine Learning Research* 4:971–1000.
- Hutter, M. 2005. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Berlin: Springer.
- Hutter, M. 2009a. Feature Dynamic Bayesian Networks. In *Proc. 2nd Conf. on Artificial General Intelligence (AGI'09)*, volume 8, 67–73. Atlantis Press.
- Hutter, M. 2009b. Feature Reinforcement Learning: Part I: Unstructured MDPs. *Journal of Artificial General Intelligence* 1:3–24.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence* 101:99–134.
- Kearns, M., and Koller, D. 1999. Efficient Reinforcement Learning in Factored MDPs. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, 740–747. San Francisco: Morgan Kaufmann.
- Koller, D., and Parr, R. 1999. Computing Factored Value Functions for Policies in Structured MDPs,. In *Proc. 16th International Joint Conf. on Artificial Intelligence (IJCAI'99)*, 1332–1339.
- Koller, D., and Parr, R. 2000. Policy Iteration for Factored MDPs. In *Proc. 16th Conference on Uncertainty in Artificial Intelligence (UAI-00)*, 326–334. San Francisco, CA: Morgan Kaufmann.
- Legg, S., and Hutter, M. 2007. Universal Intelligence: A Definition of Machine Intelligence. *Minds & Machines* 17(4):391–444.
- Lewis, D. D. 1998. Naive (Bayes) at Forty: The Independence Assumption in Information Retrieval. In *Proc. 10th European Conference on Machine Learning (ECML'98)*, 4–15. Chemnitz, DE: Springer.
- Littman, M. L.; Sutton, R. S.; and Singh, S. P. 2001. Predictive Representations of State. In *Advances in Neural Information Processing Systems*, volume 14, 1555–1561. MIT Press.
- McCallum, A. K. 1996. *Reinforcement Learning with Selective Perception and Hidden State*. Ph.D. Dissertation, Department of Computer Science, University of Rochester.
- Puterman, M. L. 1994. *Markov Decision Processes — Discrete Stochastic Dynamic Programming*. New York, NY: Wiley.
- Ross, S.; Pineau, J.; Paquet, S.; and Chaib-draa, B. 2008. Online Planning Algorithms for POMDPs. *Journal of Artificial Intelligence Research* 2008(32):663–704.
- Russell, S. J., and Norvig, P. 2003. *Artificial Intelligence. A Modern Approach*. Englewood Cliffs, NJ: Prentice-Hall, 2nd edition.

- Singh, S.; Littman, M.; Jong, N.; Pardoe, D.; and Stone, P. 2003. Learning Predictive State Representations. In *Proc. 20th International Conference on Machine Learning (ICML'03)*, 712–719.
- Singh, S. P.; James, M. R.; and Rudary, M. R. 2004. Predictive State Representations: A New Theory for Modeling Dynamical Systems. In *Proc. 20th Conference in Uncertainty in Artificial Intelligence (UAI'04)*, 512–518. Banff, Canada: AUAI Press.
- Strehl, A. L.; Diuk, C.; and Littman, M. L. 2007. Efficient Structure Learning in Factored-State MDPs. In *Proc. 27th AAAI Conference on Artificial Intelligence*, 645–650. Vancouver, BC: AAAI Press.
- Sutton, R. S., and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 2nd edition.
- Szita, I., and Lörincz, A. 2008. The Many Faces of Optimism: a Unifying Approach. In *Proc. 12th International Conference (ICML 2008)*, volume 307, 1048–1055.

Appendix A. List of Notation

Interface structures

$\mathcal{O}$	= finite or infinite set of possible observations
$\mathcal{A}$	= (small) finite set of actions
$\mathcal{R}$	= $\{0,1\}$ or $[0, R_{max}]$ or other set of rewards
$n \in \mathbb{N}$	= current time
$o_t a_t r_t$	= $oar_t \in \mathcal{O} \times \mathcal{A} \times \mathcal{R}$ = true observation, action, reward at time t

Internal structures for Φ DBN

$\log$	= binary logarithm
$t \in \{1, \dots, n\}$	= any time
$i \in \{1, \dots, m\}$	= any feature index
$\mathbf{z} = z_{1:n}$	= $z_1 \dots z_n$ (any z)
$z_+, z_*, z_\bullet$	= $\sum_j z_j, \bigcup_j z_j, (z_1, \dots, z_l)$ (any z, j, l)
$\hat{X}$	= estimate of X (any X)
$\mathcal{H}$	= $(\mathcal{O} \times \mathcal{A} \times \mathcal{R})^* \times \mathcal{O}$ = possible histories
h_n	= $oar_{1:n-1} o_n$ = actual history at time n
$\mathcal{S}$	= $\{0,1\}^m$ = internal feature/state space (can vary with n)
$\mathbf{x}$	= $(x^1, \dots, x^m) \in \mathcal{S}$ = feature vector
$w_k = R_{x^i}^i$	= local reward of feature i
Pa^i	$\subseteq \{1, \dots, m\}$ = set of parents of feature i
$\mathbf{u}^i$	$\subseteq \{x^1, \dots, x^m\}$ = sublist of feature values on which x^i depends
$\Phi: \mathcal{H} \rightarrow \mathcal{S}$	= state or feature summary of history
$s_t = \mathbf{x}_t$	= $\Phi(h_t) \in \mathcal{S}$ = realized state/feature at time t
$P(\cdot)$	= probability over states/features and rewards or parts thereof
$\text{CL}(\cdot)$	= code length
MDP	= $(\mathcal{S}, \mathcal{A}, T, R)$ = Markov Decision Process
$T_{\mathbf{x}\mathbf{x}'}^a$	= $P(\mathbf{x}_t = \mathbf{x}' \mathbf{x}_{t-1} = \mathbf{x}, a_{t-1} = a)$ = transition matrix
$\mathbf{u}^i \xrightarrow{a} x^i$	= action a in features $\mathbf{u}^i$ resulted in feature x^i
$\mathcal{T}_{\mathbf{u}^i x^i}^{ia}$	= set of times $t \in \{1, \dots, n\}$ at which $\mathbf{u}^i \xrightarrow{a} x^i$
$n_{\mathbf{u}^i x^i}^{ia}$	= $ \mathcal{T}_{\mathbf{u}^i x^i}^{ia} $ number of times $t \in \{1, \dots, n\}$ at which $\mathbf{u}^i \xrightarrow{a} x^i$
$\text{Cost}(\Phi h)$	= cost (evaluation function) of Φ based on history h
$Q_{\mathbf{x}}^*, V_{\mathbf{x}}^*$	= optimal (Q) Value (of action a) in state $\mathbf{x}$
$\gamma \in [0;1)$	= discount factor ($(1-\gamma)^{-1}$ is effective horizon)