



Bounded Rational Decision-Making with Adaptive Neural Network Priors

Heinke Hihn^(✉), Sebastian Gottwald, and Daniel A. Braun

Faculty of Engineering, Computer Science, and Psychology,
Institute for Neural Information Processing, Ulm University,
Ulm, Germany

{heinke.hihn,sebastian.gottwald,daniel.braun}@uni-ulm.de

Abstract. Bounded rationality investigates utility-optimizing decision-makers with limited information-processing power. In particular, information theoretic bounded rationality models formalize resource constraints abstractly in terms of relative Shannon information, namely the Kullback-Leibler Divergence between the agents' prior and posterior policy. Between prior and posterior lies an anytime deliberation process that can be instantiated by sample-based evaluations of the utility function through Markov Chain Monte Carlo (MCMC) optimization. The most simple model assumes a fixed prior and can relate abstract information-theoretic processing costs to the number of sample evaluations. However, more advanced models would also address the question of learning, that is how the prior is adapted over time such that generated prior proposals become more efficient. In this work we investigate generative neural networks as priors that are optimized concurrently with anytime sample-based decision-making processes such as MCMC. We evaluate this approach on toy examples.

Keywords: Bounded rationality · Variational Autoencoder
Adaptive priors · Markov Chain Monte Carlo

1 Introduction

Intelligent agents are usually faced with the task of optimizing some utility function \mathbf{U} that is a priori unknown and can only be evaluated sample-wise. We do not restrict ourselves on the form of this function, thus in principle it could be a classification or regression loss, a reward function in a reinforcement learning environment or any other utility function. The framework of information-theoretic bounded rationality [16, 17] and related information-theoretic models [3, 14, 20, 21, 23] provide a formal framework to model agents that behave in a computationally restricted manner by modeling resource constraints through information-theoretic constraints. Such limitations also lead to the emergence of hierarchies and abstractions [5], which can be exploited to reduce computational and search effort. Recently, the main principles have been successfully applied

to spiking and artificial neural networks, in particular feedforward-neural network learning problems, where the information-theoretic constraint was mainly employed as some kind of regularization [7, 11, 12, 18]. In this work we introduce bounded rational decision-making with adaptive generative neural network priors. We investigate the interaction between anytime sample-based decision-making processes and concurrent improvement of prior policies through learning, where the prior policies are parameterized as Variational Autoencoders [10]—a recently proposed generative neural network model.

The paper is structured as follows. In Sect. 2 we discuss the basic concepts of information-theoretic bounded rationality, sampled-based interpretations of bounded rationality in the context of Markov Chain Monte Carlo (MCMC), and the basic concepts of Variational Autoencoders. In Sect. 3 we present the proposed decision-making model by combining sample-based decision-making with concurrent learning of priors parameterized by Variational Autoencoders. In Sect. 4 we evaluate the model with toy examples. In Sect. 5 we discuss our results.

2 Methods

2.1 Bounded Rational Decision Making

The foundational concept in decision-making theory is Maximum Expected Utility [22], whereby an agent is modeled as choosing actions such that it maximizes its expected utility

$$\max_{p(a|w)} \sum_w \rho(w) \sum_a p(a|w) \mathbf{U}(w, a), \quad (1)$$

where a is an action from the action space A and w is a world state from the world state space W , and $\mathbf{U}(w, a)$ is a utility function. We assume that the world states are distributed according to a known and fixed distribution $\rho(w)$ and that the world states w are finite and discrete. In the case of a single world state or world state distribution $\rho(w) = \delta(w - w_0)$, the decision-making problem simplifies into a single function optimization problem $a^* = \arg \max_a \mathbf{U}(a)$. In many cases, solving such optimization problems may require an exhaustive search, where simple enumeration is extremely expensive.

A bounded rational decision maker tackles the above decision-making problem by settling on a good enough solution. Finding a bounded optimal policy requires to maximize the utility function while simultaneously remaining within some given constraints. The resulting policy is a conditional probability distribution $p(a|w)$, which essentially consists of choosing an action a given a particular world state w . The constraints of limited information-processing resources can be formalized by setting an upper bound on the D_{KL} (say B bits) that the decision-maker is maximally allowed to spend to transform its prior strategy into a posterior strategy through deliberation. This results in the following constrained optimization problem [5]:

$$\max_{p(a|w)} \sum_w \rho(w) \sum_a p(a|w) \mathbf{U}(w, a), \text{ s.t. } D_{\text{KL}}(p(a|w)||p(a)) \leq B. \quad (2)$$

This constrained optimization problem can be formulated as an unconstrained problem [16]:

$$\max_{p(a|w)} \left(\sum_w \rho(w) \sum_a p(a|w) \mathbf{U}(w, a) - \frac{1}{\beta} \text{D}_{\text{KL}}(p(a|w)||p(a)) \right), \quad (3)$$

where the inverse temperature $\beta \in \mathbb{R}^+$ is a Lagrange multiplier that influences the trade off between expected utility gain and information cost. For $\beta \rightarrow \infty$ the agent behaves perfectly rational and for $\beta \rightarrow 0$ the agent can only act according to the prior policy. The optimal prior policy in this case is given by the marginal $p(a) = \sum_{w \in W} \rho(w) p(a|w)$ [5], in which case the Kullback-Leibler divergence becomes equal to the mutual information, i.e. $\text{D}_{\text{KL}}(p(a|w)||p(a)) = I(W; A)$. The solution to the optimization problem (3) can be found by iterating the following set of self-consistent equations [5]:

$$\begin{cases} p(a|w) = \frac{1}{Z(w)} p(a) \exp(\beta_1 \mathbf{U}(w, a)) \\ p(a) = \sum_w \rho(w) p(a|w), \end{cases}$$

where $Z(w) = \sum_a p(a) \exp(\beta_1 \mathbf{U}(w, a))$ is normalization factor. Computing such a normalization factor is usually computationally expensive as it involves summing over spaces with high cardinality. We avoid this by Monte Carlo approximation.

2.2 MCMC as Sample-Based Bounded Rational Decision-Making

Monte Carlo methods are mostly used to solve two related kinds of problems. One is to generate samples x from a given distribution $q(x)$ and the other is to estimate the expectation of a function. For example, if $g(x)$ is a function for which we need to compute the expectation $\hat{\Phi} = \mathbb{E}_{q(x)}[g(x)]$ we can draw N samples $\{x_i\}_{i=1}^N$ to obtain the estimate $\hat{\Phi} = \frac{1}{N} \sum_{i=1}^N g(x_i)$ [15]. Samples can be drawn by employing Markov Chains to simulate stochastic processes. A Markov Chain can be defined by an initial probability $p^0(x)$ and a transition probability $\mathbf{T}(x', x)$, which gives the probability of transitioning from state x to x' . The probability of being in state x' at the $(t + 1)$ -th iteration is given by:

$$p^{t+1}(x') = \sum_x \mathbf{T}(x', x) p^t(x). \quad (4)$$

Such a chain can be used to generate sample proposals from a desired target distribution $q(x)$, if the following prerequisites are met [15]. Firstly, the chain must be ergodic, i.e. the chain must converge to $q(x)$ independent of the initial distribution $p^0(x)$. Secondly, the desired distribution must be an invariant distribution of the chain. A distribution $q(x)$ is an invariant of $\mathbf{T}(x', x)$ if its probability vector is an eigenvector of the transition probability matrix. A sufficient, but not necessary condition to fulfill this requirement is detailed balance, i.e. the probability of going from state x to x' is the same as going from x' to x : $q(x) \mathbf{T}(x', x) = q(x') \mathbf{T}(x, x')$.

An MCMC chain can be viewed as a bounded rational decision-making process for a single context w in the sense that it performs an anytime optimization of a utility function $\mathbf{U}(a)$ with some precision γ and that it is initialized with a prior $p(a)$. The target distribution has to be chosen as $q(a) \propto e^{\gamma \mathbf{U}(a)}$ in this case. A decision is made with the last sample when the chain is stopped. The resource corresponds then to the number of steps the chain has taken to evaluate the function $\mathbf{U}(a)$. To find the transition probabilities $\mathbf{T}(x', x)$ of the chain, we assume detailed balance and a Metropolis-Hastings scheme $\mathbf{T}(x', x) = g(x'|x)A(x'|x)$ such that

$$\frac{\mathbf{T}(x', x)}{\mathbf{T}(x, x')} = \frac{g(x'|x)A(x'|x)}{g(x|x')A(x|x')} = e^{\gamma(\mathbf{U}(x') - \mathbf{U}(x))} \quad (5)$$

with a proposal distribution $g(x'|x)$ and an acceptance probability $A(x'|x)$. One common choice that satisfies Eq. (5) is

$$A(x'|x) = \min \left\{ 1, \frac{g(x'|x)}{g(x|x')} e^{\gamma(\mathbf{U}(x') - \mathbf{U}(x))} \right\}, \quad (6)$$

which can be further simplified when using a symmetric proposal distribution with $g(x'|x) = g(x|x')$, resulting in $A(x'|x) = \min \left\{ 1, e^{\gamma(\mathbf{U}(x') - \mathbf{U}(x))} \right\}$.

Note that the decision of the chain will in general follow a non-equilibrium distribution, but that we can use the bounded rational optimum as a normative baseline to quantify how efficiently resources are used by analyzing how closely the bounded rational equilibrium is approximated.

2.3 Representing Prior Strategies with Variational Autoencoders

While an anytime optimization process such as MCMC can be regarded as a transformation from prior to posterior, the question remains how to choose the prior. While the prior may be assumed to be fixed, it would be far more efficient if the prior itself were subjected to an optimization process that minimizes the overall information-processing costs. Since in the case of multiple world states w the optimal prior is given by the marginal $p(a) = \sum_w \rho(w)p(a|w)$, we can use the outputs a of the anytime decision-making process to train a generative model of the prior $p(a)$. If the generative model was chosen from a parametric family such as a Gaussian distribution, then training would consist in updating mean and variance of the Gaussian. Choosing such a parametric family imposes restrictions on the shape of the prior, in particular in the continuous domain. Therefore, we investigate non-parametric generative models of the prior, in particular neural network models such as Variational Autoencoders (VAEs).

VAEs were introduced by [10] as generative models that use a similar architecture as deterministic autoencoder networks. Their functioning is best understood as variational Bayesian inference in a latent variable model $p(x|z, \theta)$ with prior $p(z)$, where x is observable data, and z is the latent variable that explains the data, but that cannot be observed directly. The aim is to find a parameter $\hat{\theta}_{ML}$ that maximizes the likelihood of the data $p(x|\theta) = \int p(x|z, \theta)p(z)dz$. Samples

from $p(x|\theta)$ can then be generated by first sampling z and then sampling an x from $p(x|z, \theta)$. As the maximum likelihood optimization may prove difficult due to the integral, we may express the likelihood in a different form by assuming a distribution $q(z|x, \eta)$ such that

$$\begin{aligned} \log p(x|\theta, \eta) &= \int q(z|x, \eta) \log \frac{p(x|z, \theta)p(z)}{q(z|x, \eta)} dz + \underbrace{\int q(z|x, \eta) \log \frac{q(z|x, \eta)}{p(z|x, \theta)} dz}_{=D_{\text{KL}}(q||p) \geq 0} \\ &\geq \int q(z|x, \eta) \log \frac{p(x|z, \theta)p(z)}{q(z|x, \eta)} dz =: F(\theta, \eta). \end{aligned} \quad (7)$$

Assuming that the distribution $q(z|x, \eta)$ is expressive enough to approximate the true posterior $p(z|x, \theta)$ reasonably well, we can neglect the D_{KL} between the two distributions, and directly optimize the lower bound $F(\theta, \eta)$ through gradient descent. In VAEs $q(z|x, \eta)$ is called the encoder that translates from x to z and $p(x|z, \theta)$ is called the decoder that translates from z to x . Both distributions and the prior $p(z)$ are assumed to be Gaussian

$$\begin{aligned} p(x|z, \theta) &= \mathcal{N}(x|\mu_\theta(z), \sigma^2 \mathbb{I}) \\ q(z|x, \eta) &= \mathcal{N}(z|\mu_\eta(x), \Sigma_\eta(x)) \\ p(z) &= \mathcal{N}(z|0, \mathbb{I}), \end{aligned}$$

where $\mu_\theta(z)$, $\mu_\eta(x)$ and $\Sigma_\eta(x)$ are non-linear functions implemented by feed-forward neural networks and where it is ensured that $\sigma^2 \searrow 0$ and that $\Sigma_\eta(x)$ is a covariance matrix.

Note that the optimization of the autoencoder itself can also be viewed as a bounded rational choice

$$\max_{\theta, \eta} \left(\mathbb{E}_{q(z|x, \eta)} [\log p(x|z, \theta)] - D_{\text{KL}}(q(z|x, \eta)||p(z)) \right), \quad (8)$$

where the expected likelihood is maximized while the encoder distribution $q(z|x, \eta)$ is kept close to the prior $p(z)$.

3 Modeling Bounded Rationality with Adaptive Neural Network Priors

In this section we combine MCMC anytime decision-processes with adaptive autoencoder priors. In the case of a single world state, the combination is straightforward in that each decision selected by the MCMC process is fed as an observable input to an autoencoder. The updated autoencoder is then used as an improved prior to initialize the next MCMC decision. In case of multiple world states, there are two straightforward scenarios. In the first scenario there are as many priors as world states and each of them is updated independently. For each world state we obtain exactly the same solution as in the single world

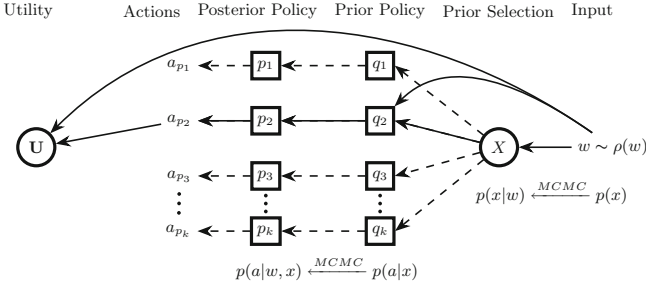


Fig. 1. For each incoming world state w our model samples a prior indexed by $x_i \sim p(x|w)$. Each prior $p(a|x)$ is represented by a VAE. To arrive at the posterior policy $p(a|w, x)$, an anytime MCMC optimization is seeded with $a_0 \sim p(a|x)$ to generate a sample from $p(a|w, x)$. The prior selection policy is also implemented by an MCMC chain and selects agents that have achieved high utility on a particular w .

state case. In the second scenario there is only a single prior over actions for all world states. In this case the autoencoder is trained with the decisions by all MCMC chains such that the autoencoder should converge to the optimal rate distortion prior. A third, more interesting scenario occurs when we allow multiple priors, but less than world states—compare Fig. 1. This is especially plausible when dealing with continuous world states, but also in the case of large discrete spaces.

3.1 Decision Making with Multiple Priors

Decision-making with multiple priors can be regarded as a multi-agent decision-making problem where several bounded rational decision-makers are combined into a single decision-making process [5]. In our case the most suitable arrangement of decision-makers is a two-step process where first each world state is assigned probabilistically to a prior which is then used in the second step to initialize an MCMC chain—compare Fig. 1. The output of that chain is then used to train the autoencoder corresponding to the selected prior. As each prior may be responsible for multiple world states, each prior will learn an abstraction that is specialized for this subspace of world states. This two-stage decision-process can be formalized as a bounded rational optimization problem

$$\max_{p(a|w,x), p(x|w)} \left(\mathbb{E}_{p(a|w,x)}[\mathbf{U}(w, a)] - \frac{1}{\beta_1} I(W; X) - \frac{1}{\beta_2} I(W; A|X) \right), \quad (9)$$

where $p(x|w)$ is selecting the responsible prior $p(a|x)$ indexed by x for world state w . The resource parameter for the first selection stage is given by β_1 and by β_2 for the second decision made by the MCMC process. The solution of optimization (9) is given by the following set of equations:

$$\left\{ \begin{array}{l} p(x|w) = \frac{1}{Z(w)} p(x) \exp(\beta_1 \Delta F_{\text{par}}(w, x)) \\ p(x) = \sum_w \rho(w) p(x|w) \\ p(a|w, x) = \frac{1}{Z(w, x)} p(a|x) \exp(\beta_2 \mathbf{U}(w, a)) \\ p(a|x) = \sum_w p(w|x) p(a|w, x) \\ \Delta F_{\text{par}}(w, x) = \mathbb{E}_{p(a|w, x)}[\mathbf{U}(w, a)] - \frac{1}{\beta_2} \text{D}_{\text{KL}}(p(a|w, x) || p(a|x)), \end{array} \right. \quad (10)$$

where $Z(w)$ and $Z(w, x)$ are the normalization factors and $\Delta F_{\text{par}}(w, x)$ is the free energy of the action selection stage. The marginal distribution $p(a|x)$ encapsulates an action selection policy consisting of the priors $p(a|w, x)$ weighted by the responsibilities given by the Bayesian posterior $p(w|x)$. Note that the Bayesian posterior is not determined by a given likelihood model, but is the result of the optimization process (9).

3.2 Model Architecture

Equation (10) describe abstractly how a two-step decision process with bounded rational decision-makers should be optimally partitioned. In this section we propose a sample-based model of a bounded rational decision process that approximately corresponds to Eq. (10) such that the performance of the decision process can be compared against its normative baseline. To translate Eq. (10) into a stochastic process we proceed in three steps. First, we implement the priors $p(a|x)$ as Variational Autoencoders. Second, we formulate an MCMC chain that is initialized with a sample from the prior and generates a decision $a \sim p(a|x, w)$. Third, we design an MCMC chain that functions as a selector between the different priors.

Autoencoder Priors. Each prior $p(a|x)$ in Eq. (10) is represented by a VAE that learns to generate action samples that mimic the samples given by the MCMC chains—compare Fig. 2. The functions $\mu_\theta(z)$, $\mu_\eta(a)$ and $\Sigma_\eta(a)$ are implemented as feed-forward neural networks with one hidden layer. The units in the hidden layer were all chosen with sigmoid activation function, the output units in the case of the μ -functions were also chosen as sigmoids and for the Σ -function as ReLU. During training the weights η and θ are adapted to optimize the expected log-likelihood of the action samples that are given by the decisions made by the MCMC chains for all world states that have been assigned to the prior $p(a|x)$. Due to the Gaussian shape of the decoder distribution, optimizing the log-likelihood corresponds to minimizing quadratic loss of the reconstruction error. After training, the network can generate sample actions itself by feeding the decoder network with samples from $\mathcal{N}(z|0, \mathbb{I})$.

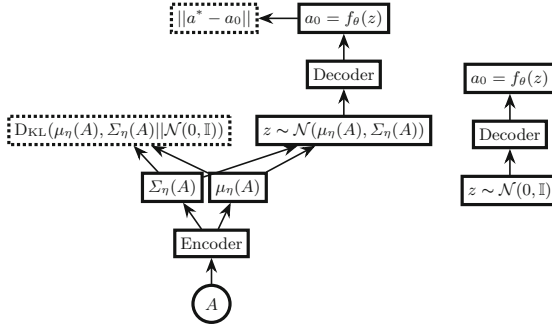


Fig. 2. The encoder translates the observed action into a latent variable z , whereas the decoder translates the latent variable z into a proposed action a . During training the weights η and θ are adapted to optimize the expected log-likelihood of the observed action samples. After training, the network can generate actions by feeding the decoder network with samples from $\mathcal{N}(z|0, \mathbb{I})$.

MCMC Decision-Making. To implement the bounded rational decision-maker $p(a|w, x)$ we obtain an action sample $a \sim p(a|x)$ from the autoencoder prior to initialize an MCMC chain that optimizes the target utility $\mathbf{U}(w, a)$ for the given world state. We run the MCMC chain for N_{\max} steps. In each step we generate a proposal from a Gaussian distribution with $g(a'|a) = \mathcal{N}(a'|a, \sigma^2)$ and accept with probability

$$A(a'|a) = \min \{1, \exp(\gamma(\mathbf{U}(w, a') - \mathbf{U}(w, a)))\}. \quad (11)$$

Over the course of N_{\max} time steps, the precision γ is adjusted following an annealing schedule conditioned on the maximum number of steps N_{\max} . We use an inverse Boltzmann annealing schedule, i.e. $\gamma^{(k)} = \gamma^0 + \alpha \log(1 + k)$, where α is a tuning parameter. The rationale behind this is that we assume the sampling process to be coarse grained in the beginning and is getting finer during the search.

Prior Selection. To implement the bounded rational prior selection $p(x|w)$ through an MCMC process, we first sample an x from the prior $p(x)$ and start an MCMC chain that (approximately) optimizes $\Delta F_{\text{par}}(w, x)$ for a given world state w sampled from $\rho(w)$. The prior $p(x)$ is represented by a multinomial and updated by the frequencies of the selected prior indices x . The number of steps in the prior selection MCMC chain was kept constant at a value of N_{\max}^{sel} and similarly the precision γ^{sel} was annealed over the course of N_{\max}^{sel} time steps. The target $\Delta F_{\text{par}}(w, x)$ comprises a trade-off between expected utility and information resources. However, it cannot be directly evaluated and would require the computation of $D_{\text{KL}}(p(a|x, w) \| p(a|x))$. Here we use number of steps in the downstream MCMC process as a resource measure. As the number of

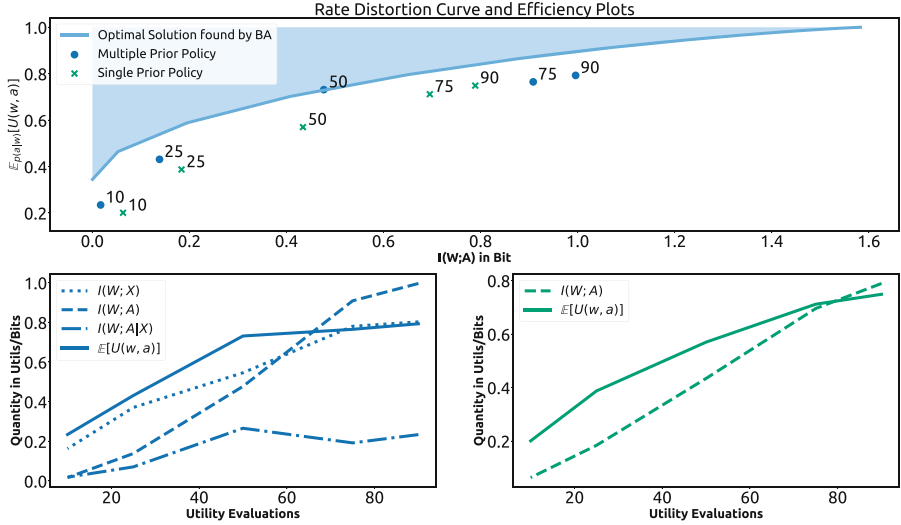


Fig. 3. Top: The line is given by the Rate Distortion Curve that forms a theoretical efficiency frontier, characterized by the ratio between mutual information and expected utility. Crosses represent single-prior agents and dots multi-prior systems. The labels indicate how many steps were assigned to the second MCMC chain of a total of 100 steps. Bottom: Information processing and expected utility is increasing in the number of utility evaluations, as we expected.

downstream steps was constant, the model selector’s choice only depended on the average utility achieved by each decision-maker, which results in the acceptance rule

$$A(x'|x) = \min \left\{ 1, \exp(\gamma^{\text{sel}} (\mathbb{E}_{p(a|w,x)}[\mathbf{U}(w, a)] - \mathbb{E}_{p(a|w,x')}[\mathbf{U}(w, a)])) \right\}.$$

As the priors are discrete choices the proposal distribution $q(x_p|x_p)$ samples globally with $p(x) = \frac{1}{|X|}$ for all x .

4 Empirical Results

To demonstrate our approach we evaluate two scenarios. First, a simple agent, which is equipped with a single prior policy $p_\eta(a)$, as introduced in Sect. 2. In case of a single agent there is no need for a prior selection stage. Second, we evaluated a multi-prior decision-making system and compared the results to the single prior agent. For the multi-prior agent, we split a fixed number of MCMC steps between the prior selection and the action selection. The task we designed consists of six world states where each world state has a Gaussian utility function in the interval $[0, 1]$ with a unique optimum. In both settings, we equipped the Variational Autoencoders with one hidden layer consisting of 16 units with ReLU activations. We implemented the experiments using Keras [2]. We show the results in Fig. 3.

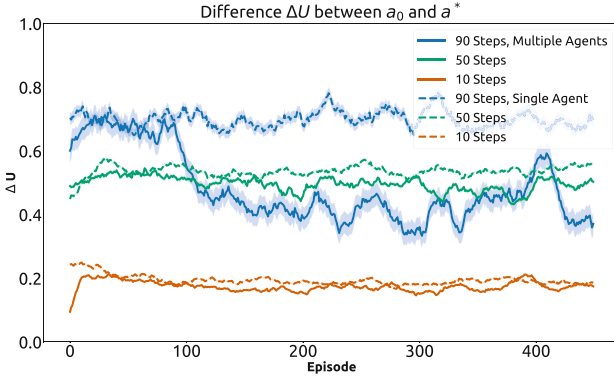


Fig. 4. Our results indicate that having multiple priors is more beneficial, if more steps are available in total. Note that the stochasticity of our method decreases with the number of allowed steps, as shown by the uncertainty band (transparent regions).

Our results indicate that using MCMC evaluation steps as a surrogate for information processing costs can be interpreted as bounded rational decision-making. In Fig. 3 we show the efficiency of several agents with different processing constraints. To compare our results to the theoretical baseline, we discretized the action space into 100 equidistant slices and solved the problem using the algorithm proposed in [5] to implement Eq. (10). Furthermore our results indicate that the multi-prior system generally outperforms the single-prior system in terms of utility.

To illustrate the differences in efficiency between the single prior agent and the multi-prior agents, we plotted in Fig. 4 utility gained through the second MCMC optimization. For multi-prior agents this is caused by specialized priors which provide initializations to the MCMC chains that are close to the optimal action. In this particular case, ΔU does not become zero because we allow only three priors to cover six world states, thus leading to abstraction, i.e. specializing on actions that fit well for the assigned world states. In single-prior agents, the prior is adapting to all world states, thus providing, on average, an initial action that is suboptimal for the requested world state.

5 Discussion

In this study we implemented bounded rational decision makers with adaptive priors. We achieved this with Variational Autoencoder priors. The bounded rational decision-making process was implemented by MCMC optimization to find the optimal posterior strategy, thus giving a computationally simple way of generating samples. As the number of steps in the optimization process was constrained, we could quantify the information processing capabilities of the resulting decision-makers using relative Shannon entropy. Our analysis may have interesting implications, as it provides a normative framework for this kind of

combined optimization of adaptive priors and decision-making processes. Prior to our work there have been several attempts to apply the framework of information-theoretic bounded rationality to machine learning tasks [7, 11, 12, 18]. The novelty of our approach is that we design adaptive priors for both the single-step case and the multi-agent case and we demonstrate how to transform information-theoretic constraints into computational constraints in the form of MCMC steps.

Recently, the combination of Monte Carlo optimization and neural networks has gained increasing popularity. These approaches include both using MCMC processes to find optimal weights in ANNs [1, 4] and using ANNs as parametrized proposal distributions in MCMC processes [8, 13]. While our approach is more similar to the latter, the important difference is that in such adaptive MCMC approaches there is only a single MCMC chain with a single (adaptive) proposal to optimize a single task, whereas in our case there are multiple adaptive priors to initialize multiple chains with otherwise fixed proposal, which can be used to learn multiple tasks simultaneously. In that sense our work is more related to mixture-of-experts methods and divide-and-conquer paradigms [6, 9, 24], where we employ a selection policy rather than a blending policy, as we design our model specifically to encourage specialization. In mixture-of-experts models, there are multiple decision-makers that correspond to multiple priors in our case, but experts are typically not modeled as anytime optimization processes. The possibly most popular combination of neural network learning with Monte Carlo methods was achieved by AlphaGo [19], which beat the leading Go champion by optimizing the strategies provided by value networks and policy networks with Monte Carlo Tree Search, leading to a major breakthrough in reinforcement learning. An important difference here is that the neural network is used to directly approximate the posterior and MCMC is used to improve performance by concentrating on the most promising moves during learning, whereas in our case ANNs are used to represent the prior. Moreover, in our work we assumed the utility function (i.e. the value network) to be given. For future work it would be interesting to investigate how to incorporate learning the utility function into our model to investigate more complex scenarios such as in reinforcement learning.

Acknowledgement. This work was supported by the European Research Council Starting Grant *BRISC*, ERC-STG-2015, Project ID 678082.

References

1. Andrieu, C., De Freitas, N., Doucet, A.: Reversible jump MCMC simulated annealing for neural networks. In: Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence, pp. 11–18. Morgan Kaufmann Publishers Inc. (2000)
2. Chollet, F., et al.: Keras (2015). <https://keras.io>
3. Vul, E., Goodman, N., Griffiths, T.L., Tenenbaum, J.B.: One and done? Optimal decisions from very few samples. *Cogn. Sci.* **38**(4), 599–637 (2014)

4. Freitas, J., Niranjan, M., Gee, A.H., Doucet, A.: Sequential Monte Carlo methods to train neural network models. *Neural Comput.* **12**(4), 955–993 (2000)
5. Genewein, T., Leibfried, F., Grau-Moya, J., Braun, D.A.: Bounded rationality, abstraction, and hierarchical decision-making: an information-theoretic optimality principle. *Front. Robot. AI* **2**, 27 (2015)
6. Ghosh, D., Singh, A., Rajeswaran, A., Kumar, V., Levine, S.: Divide-and-conquer reinforcement learning. arXiv preprint [arXiv:1711.09874](https://arxiv.org/abs/1711.09874) (2017)
7. Grau-Moya, J., Leibfried, F., Genewein, T., Braun, D.A.: Planning with information-processing constraints and model uncertainty in Markov decision processes. In: Frasconi, P., Landwehr, N., Manco, G., Vreeken, J. (eds.) *ECML PKDD 2016. LNCS (LNAI)*, vol. 9852, pp. 475–491. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46227-1_30
8. Gu, S., Ghahramani, Z., Turner, R.E.: Neural adaptive sequential Monte Carlo. In: *Advances in Neural Information Processing Systems*, pp. 2629–2637 (2015)
9. Haruno, M., Wolpert, D.M., Kawato, M.: Mosaic model for sensorimotor learning and control. *Neural Comput.* **13**(10), 2201–2220 (2001)
10. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
11. Leibfried, F., Braun, D.A.: A reward-maximizing spiking neuron as a bounded rational decision maker. *Neural Comput.* **27**(8), 1686–1720 (2015)
12. Leibfried, F., Grau-Moya, J., Ammar, H.B.: An information-theoretic optimality principle for deep reinforcement learning. arXiv preprint [arXiv:1708.01867](https://arxiv.org/abs/1708.01867) (2017)
13. Levy, D., Hoffman, M.D., Sohl-Dickstein, J.: Generalizing Hamiltonian Monte Carlo with neural networks. In: *International Conference on Learning Representations* (2018)
14. Lewis, R.L., Howes, A., Singh, S.: Computational rationality: linking mechanism and behavior through bounded utility maximization. *Top. Cogn. Sci.* **6**(2), 279–311 (2014)
15. MacKay, D.J.C.: Introduction to Monte Carlo methods. In: Jordan, M.I. (ed.) *Learning in Graphical Models. ASID*, vol. 89, pp. 175–204. Springer, Dordrecht (1998). https://doi.org/10.1007/978-94-011-5014-9_7
16. Ortega, P.A., Braun, D.A.: Thermodynamics as a theory of decision-making with information-processing costs. *Proc. R. Soc. Lond. A: Math. Phys. Eng. Sci.* **469**(2153) (2013)
17. Ortega, P.A., Braun, D.A., Dyer, J., Kim, K.E., Tishby, N.: Information-theoretic bounded rationality. arXiv preprint [arXiv:1512.06789](https://arxiv.org/abs/1512.06789) (2015)
18. Peng, Z., Genewein, T., Leibfried, F., Braun, D.A.: An information-theoretic online update principle for perception-action coupling. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 789–796. IEEE (2017)
19. Silver, D., et al.: Mastering the game of go with deep neural networks and tree search. *Nature* **529**(7587), 484–489 (2016)
20. Tishby, N., Polani, D.: Information theory of decisions and actions. In: Cutsuridis, V., Hussain, A., Taylor, J. (eds.) *Perception-Action Cycle: Models, Architectures, and Hardware. SSCNS*, pp. 601–636. Springer, New York (2011). https://doi.org/10.1007/978-1-4419-1452-1_19
21. Todorov, E.: Efficient computation of optimal actions. *Proc. Natl. Acad. Sci.* **106**(28), 11478–11483 (2009)
22. Von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behavior*, Commemorative edn. Princeton University Press, Princeton (2007)

23. Wolpert, D.H.: Information theory - the bridge connecting bounded rational game theory and statistical physics. In: Braha, D., Minai, A., Bar-Yam, Y. (eds.) *Complex Engineered Systems: Science Meets Technology*. UCS, pp. 262–290. Springer, Heidelberg (2006). https://doi.org/10.1007/3-540-32834-3_12
24. Yuksel, S.E., Wilson, J.N., Gader, P.D.: Twenty years of mixture of experts. *IEEE Trans. Neural Netw. Learn. Syst.* **23**(8), 1177–1193 (2012)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

