

# Computing Complexity-aware Plans Using Kolmogorov Complexity

Elis Stefansson<sup>1</sup> and Karl H. Johansson<sup>1</sup>

**Abstract**—In this paper, we introduce complexity-aware planning for finite-horizon deterministic finite automata with rewards as outputs, based on Kolmogorov complexity. Kolmogorov complexity is considered since it can detect computational regularities of deterministic optimal policies. We present a planning objective yielding an explicit trade-off between a policy's performance and complexity. It is proven that maximising this objective is non-trivial in the sense that dynamic programming is infeasible. We present two algorithms obtaining low-complexity policies, where the first algorithm obtains a low-complexity optimal policy, and the second algorithm finds a policy maximising performance while maintaining local (stage-wise) complexity constraints. We evaluate the algorithms on a simple navigation task for a mobile robot, where our algorithms yield low-complexity policies that concur with intuition.

## I. INTRODUCTION

### A. Motivation

Artificial intelligence has under the last decade progressed significantly achieving superhuman performance in challenging domains such as the video game Atari and the board game Go [22], [31]. Unfortunately, for more complex and unconstrained environments (e.g., advanced real-world systems such as autonomous vehicles) results are more limited [25]. One major challenge here is the huge space of all possible strategies (policies) that the agent (i.e., robot or machine) can perform, making tractable solutions cumbersome to obtain naively. However, humans tend to face these complex domains with relative ease.

One explanation why humans perform well in complex tasks comes from cognitive neuroscience, proposing that general intelligence is linked to efficient compression, known as the efficient coding hypothesis [3], [4], [32]. This idea is not new but can be traced back to William of Occam saying “*If there are alternative explanations for a phenomenon, then, all other things being equal, we should select the simplest one*”, where the simple alternative is the alternative with the shortest explanation. This methodology is known as Occam's razor [19]. Formalisations of Occam's razor have been able to detect computational regularities in for example the decimals of  $\pi$  [40] and automatically extract low-complexity physical laws (e.g.,  $E = mv^2/2$ ) directly from data [37]. In this paper, we are interested if a similar formalisation can automatically extract low-complexity policies, seen as a first step towards more tractable and intelligent behaviour for agents acting in complex environments.

<sup>1</sup>School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Sweden. Email: {elisst, kallej}@kth.se

This work was partially funded by the Swedish Foundation for Strategic Research, the Swedish Research Council, and the Knut och Alice Wallenberg foundation.

### B. Contribution

The main contribution of this paper is to define a complexity measure for policies in deterministic finite automata (DFA) [14] using Kolmogorov complexity [19], and to construct tractable complexity-aware planning algorithms with explicit trade-offs between performance and complexity based on this measure. More precisely, our contributions are three-fold:

Firstly, we define a complexity measure for deterministic policies in finite-horizon DFA with rewards as outputs. This measure uses Kolmogorov complexity to evaluate how complex a policy is to execute. Kolmogorov complexity, which can be seen as a formalisation of Occam's razor, is a computational notion of complexity being able to not only detect statistical regularities (commonly exploited in standard information theory) but also computational regularities.<sup>1</sup> Our key insight is that optimal policies typically possess computational regularities (such as reaching a goal state) apart from statistical ones, making Kolmogorov complexity an appealing complexity evaluator.

Secondly, we present a complexity-aware planning objective based on our complexity measure, yielding an explicit trade-off between a policy's performance and complexity. We also prove that maximising this objective is non-trivial in the sense that dynamic programming [5] is infeasible.

Thirdly, we present two algorithms obtaining low-complexity policies. The first algorithm, Complexity-guided Optimal Policy Search (COPS), finds a policy with low complexity among all optimal policies, following a two-step procedure. In the first step, the algorithm finds all optimal policies, without any complexity constraints. Dropping the complexity constraints, this step can be done using dynamic programming. The second step runs a uniform-cost search [29] over all optimal policies, guided by a complexity-driven cost, favouring low-complexity policies and enabling moderate search depths. The second algorithm, Stage-complexity-aware Program (SCAP), penalises instead policies locally for executing complex manoeuvres. This is done by partitioning the horizon into stages, with local complexity constraints over the stages instead of the full horizon, which enables dynamic programming over the stages. Finally, we evaluate our algorithms on a simple navigation task where a mobile robot tries to reach a certain goal state. Our algorithms yield low-complexity policies that concur with intuition.

<sup>1</sup>An example of a sequence with computational regularity but no apparent statistical regularity is the infinite sequence 1234567891011...

### C. Related Work

The interest of complexity in control and learning has a long history going back to Bellman's curse of dimensionality [5], Witsenhausen's counterexample [38], [39] and the general open problem under what conditions LQG admits an optimal low-dimensional feedback controller [41], just to mention a few, and has lately been reviewed as an essential component for intelligent behaviour [28]. Recent attempts to find low-complexity policies can be divided into two categories. In the first category, the system itself is approximated by a low-complexity system (e.g., smaller dimension), whereas an approximately optimal solution can be obtained. Methods in this category include bisimulation [13], [6], PCA analysis [26], [20], and information-theoretic compression such as the information bottleneck method [1], [17]. In the second category, a low-complexity policy is instead obtained directly. Here, notable methods include policy distillation [30], VC-dimension constraints [16], concise finite-state machine plans [23], [24], low-memory policies through sparsity constraints [7], and information-theoretic approaches such as KL-regularisation [27], [36], mutual information regularisation with variations [34], [12], [35], and minimal specification complexity [11], [10]. Our work belongs to this second category and resembles [23], [24], [11], [10] the most, but differ since we consider Kolmogorov complexity.

Kolmogorov complexity has also been considered in the context of reinforcement learning as a tool for complexity-constrained inference [9], [2], [15] based on Solomonoff's theory of inductive inference [19]. We differ by focusing instead on constraining the computational complexity of the obtained policy itself, assuming the underlying system to be known.

Finally, the work [8] considers Kolmogorov complexity to measure the complexity of an action sequence similar to this line of work. We differ by also optimising over the complexity, while [8] only evaluates the complexity of an immutable object.

### D. Outline

The remaining paper is as follows. Section II provides preliminaries together with the problem statement. Section III defines the complexity measure together with the complexity-aware planning objective, and proves that dynamic programming is infeasible. Section IV presents two algorithms for yielding low-complexity policies and Section V evaluates the algorithms on a mobile robot example. Finally, Section VI concludes the paper.

An extended version of this paper can be found at [33].

## II. PROBLEM FORMULATION

### A. Turing Machines

We give a brief review of Turing machines following [19]. A Turing machine is a mathematical model of a computing device, manipulating symbols on an infinite list of cells (known as the tape), reading one symbol at a time on the tape with a one access pointer (known as the head). Formally:

*Definition 1:* A Turing machine  $M$  is a tuple  $(A, Q, \delta, q_0)$  with: A finite set of tape symbols  $A \cup \{b\}$  with alphabet  $A$  and blank symbol  $b$ ; A finite set of states  $Q$  with start state  $q_0 \in Q$ ; A partial function<sup>2</sup>  $\delta : Q \times (A \cup \{b\}) \rightarrow (A \cup \{b\} \cup \{L, R\}) \times Q$  called the transition function. An execution of a Turing machine  $M$  with input string<sup>3</sup>  $x \in A^*$  starts with  $x$  on the tape (and blank symbols on both sides of  $x$ ),  $q = q_0$  as state, and head on the first symbol  $s$  of  $x$ . It then transitions according to  $(h, p) = \delta(q, s)$ , where  $h \in A \cup \{b\}$  specifies what current scanned symbol  $s$  should be replaced with, or if head should move left (right) one step ( $h \in \{L, R\}$ ), and  $p \in Q$  specifies the next state. This transition procedure is repeated until  $\delta(q, s)$  becomes undefined. In this case,  $M$  halts with output  $y \in A^*$  equal to the maximal string in  $A^*$  currently under scan, or  $y = 0$  if  $b$  is scanned. For each Turing machine  $M$ , this input-output convention defines a partial function  $\phi$  (defined when  $M$  halts) known as a *partial computable function*.

An important special class of Turing machines is the universal Turing machines. Informally, a universal Turing machine is a Turing machine that can simulate any other Turing machine, and can therefore be seen as an idealised version of a modern computer. A fundamental result states that there exist such machines [19], a fact utilised when defining the Kolmogorov complexity.

### B. Kolmogorov Complexity

In this section, we present needed theory concerning Kolmogorov complexity [19]. Informally, the Kolmogorov complexity of an object  $x$  is the length of the smallest program  $p$  on a computer that outputs  $x$ . If  $p$  is much smaller than  $x$ , then we have compressed the information in  $x$  to only the essential information of  $x$ . The Kolmogorov complexity can therefore be seen as a formalisation of Occam's razor, stripping away all the non-essential information. More formally, the computer is a universal Turing machine  $U$ , and the object  $x \in A^*$  and the program  $p \in A^*$  are strings of a finite alphabet  $A$ . Towards a precise definition, we need the following notion:

*Definition 2 ([19]):* The complexity of a string  $x \in A^*$  with respect to a partial computable function  $f : A^* \rightarrow A^*$  is defined as

$$K_f(x) = \min\{\ell(p) : f(p) = x\},$$

where  $\ell(p)$  is the length of string  $p$ . Here,  $p$  is interpreted as the program input to  $f$ , and  $K_f(x)$  is thus the length of the smallest program that, via  $f$ , describes  $x$ .

The following fundamental result, known as the invariance theorem, asserts that there is a partial computable function with shorter (i.e., more compact) descriptions than any other partial computable function, up to a constant:

*Theorem 1 (Invariance theorem [19]):* There exists a partial computable function  $\phi$ , constructed from a universal

<sup>2</sup>We use the notation  $f : A \rightarrow B$  to denote a partial function from a set  $A$  to a set  $B$  (i.e., a function that is only defined on a subset of  $A$ ).

<sup>3</sup>Here,  $A^*$  denotes the set of all finite strings from the alphabet  $A$ , e.g.,  $x = 010 \in A^*$  if  $A = \{0, 1\}$ .

Turing machine  $U$ , with the following *additively optimal* property: For any other partial computable function  $\psi$ , there exists a constant  $c_\psi$  (dependent only on  $\psi$ ) such that for all  $x \in A^*$ :  $K_\phi(x) \leq K_\psi(x) + c_\psi$ .

This function  $\phi$  serves as our computer when defining the Kolmogorov complexity:

*Definition 3 (Kolmogorov Complexity [19]):* Let  $\phi$  be as in Theorem 1. The Kolmogorov complexity of  $x \in A^*$  is  $K_\phi(x)$ . We sometimes abbreviate  $K_\phi$  as  $K$ .

The Kolmogorov complexity is robust in the sense that it depends only benignly on the choice of  $\phi$ , see [19] for details. The Kolmogorov complexity  $K(x)$  is however not computable in general, but can only be over-approximated pointwise [19]. The method we use in this paper to approximate  $K(x)$  is from [40] and based on algorithmic probability [19], see [33] for a summary. This estimation method is used in the simulations (Section V), while the underlying theoretical framework we develop (Section III and IV) is based on the exact Kolmogorov complexity.<sup>4</sup>

### C. Deterministic Finite Automata

We consider finite-horizon planning for discrete systems formalised as time-varying DFA [14] with actions as inputs and rewards as outputs (i.e., time-varying Mealy machines [21]), on the form:<sup>5</sup>

$$\mathcal{M} = \langle S, A, T, (f_t, r_t, F_t, t \in \mathbb{T}) \rangle, \quad (1)$$

where  $S$  is the finite set of states,  $A$  the finite action set,  $T \in \mathbb{N}$  the horizon length and  $\mathbb{T} = \{0, 1, \dots, T\}$ ,  $f_t : S \times A \rightarrow S$  the transition function specifying the next state  $f_t(s, a) \in S$  at time  $t$  given current state  $s \in S$  and action  $a \in A$ , and  $r_t : S \times A \rightarrow \mathbb{R}$  the corresponding received reward  $r_t(s, a) \in \mathbb{R}$ . The system stops at  $t = T$ , given by final state sets  $F_T = S$  and  $F_t = \emptyset$  for  $t < T$ . Given a start state  $s_0 \in S$ , the objective is to find a (deterministic) policy  $\pi : \mathbb{T} \times S \rightarrow A$  maximising the total reward  $\sum_{t=0}^T r_t(s_t, a_t)$  subject to the transition dynamics  $s_{t+1} = f_t(s_t, \pi(t, s_t))$ . A policy which does this for every start state  $s_0 \in S$  is called an *optimal* policy.

### D. Problem Statement

We now formalise the problem statement. In this paper, we answer the following questions:

- 1) Given an  $\mathcal{M}$  as in (1), how can one define a complexity measure  $C(s_0, \pi)$  capturing how complex a policy  $\pi$  is to execute from a start state  $s_0 \in S$ ?
- 2) How can one construct a planning objective  $\Psi$  with a formal trade-off between total reward  $\sum_{t=0}^T r_t(s_t, a_t)$  and complexity  $C(s_0, \pi)$  for a policy  $\pi$ ?
- 3) How can one construct algorithms that maximises  $\Psi$ ?

<sup>4</sup>We stress that other estimation methods can also be used, e.g., Lempel-Ziv compression [18] (e.g., used in [8]), since the theoretical framework is independent on the particular estimation method chosen. We picked [40] due to its more direct connection with Kolmogorov complexity, whereas Lempel-Ziv compression relies on classical information theory.

<sup>5</sup>The time-varying feature of the DFA can be lifted by including the time into the state. We keep the current notation for easier readability.

- 4) In particular, can dynamic programming be used to maximise  $\Psi$ ?

Questions 1, 2 and 4 are answered in Section III, while Section IV answers question 3 together with the numerical evaluations in Section V.

## III. COMPLEXITY-AWARE PLANNING

This section introduces a complexity measure in Section III-A and then sets up an appropriate complexity-aware planning objective in Section III-B. Finally, we prove that dynamic programming cannot be used to maximise this objective in Section III-C. Throughout this section, we fix a system  $\mathcal{M}$  as in (1).

### A. Execution Complexity

We start by defining our complexity measure for policies, capturing how complex it is to execute a policy  $\pi$  from a start state  $s_0 \in S$ . Towards this, note that, given a start state  $s_0 \in S$  and a policy  $\pi$ , we get a sequence of actions  $(\pi(0, s_0), \dots, \pi(T, s_T))$  in  $A^{T+1}$  from time  $t = 0$  to time  $t = T$ , generated by  $\pi$  and the dynamics  $s_{t+1} = f_t(s_t, \pi(t, s_t))$ . We denote this action sequence by  $E(s_0, \pi)$  and say that  $\pi$  has low *execution complexity* if  $K(E(s_0, \pi))$  is low. Intuitively, a policy  $\pi$  with low execution complexity has a small program that can execute it.

*Definition 4 (Execution Complexity):* Given a start state  $s_0 \in S$ , the execution complexity of a policy  $\pi$  is  $C(s_0, \pi) = K(E(s_0, \pi))$ .

### B. Complexity-aware Planning Objective

The execution complexity can be used to find policies with high total reward while keeping a low complexity. Formally, we want to find an action sequence  $a_{0:T} = (a_0, a_1, \dots, a_T)$  maximising the objective<sup>6</sup>

$$\max_{a_{0:T}} \left[ \sum_{t=0}^T r_t(s_t, a_t) - \beta K(a_{0:T}) \right], \quad (2)$$

subject to the dynamics  $s_{t+1} = f_t(s_t, a_t)$  and start state  $s_0 \in S$ . Here,  $\beta \geq 0$  determines how much we penalise complexity relative to obtaining a high total reward.

*Example 1 (Simple optimal policies):* For  $\beta > 0$  sufficiently low, we obtain, for a start state  $s_0 \in S$ , an interesting subset of all optimal policies (i.e., policies maximising the total reward  $\sum_{t=0}^T r_t(s_t, a_t)$ ) with the lowest complexity, which we here call *simple optimal policies*. Note that, the simple optimal policies can also be found via the objective

$$\min_{a_{0:T}} K(a_{0:T}) \quad \text{s.t.} \quad a_{0:T} \in \arg \max_{a_{0:T}} \sum_{t=0}^T r_t(s_t, a_t), \quad (3)$$

subject to the dynamics  $s_{t+1} = f_t(s_t, a_t)$  and start state  $s_0$ . See [33] for a proof.

<sup>6</sup>For convenience, we maximise directly over control inputs  $a_t = \pi(t, s_t)$  instead of policies, since the horizon  $T$  is typically lower than the number of states in  $S$ . With this convention,  $K(a_{0:T})$  in (2) agrees with  $C(s_0, \pi)$  in Definition 4.

### C. Dynamic Programming is Infeasible

We are interested in algorithms obtaining the maximum of (2). A standard method for finding an optimal action sequence maximising the total reward  $\sum_{t=0}^T r_t(s_t, a_t)$  is dynamic programming [5]. Hence, one may ask if (2) can be solved using dynamic programming. The answer is negative and follows from the following result.

*Proposition 1:* For sufficiently large  $T \in \mathbb{N}$ , we cannot decompose  $K$  on the form

$$K(x_{1:T}) = h_1(x_1) + g_{2:T}(x_{2:T})$$

where  $\{g_{k:T}\}_{k=2}^T$  are functions given by

$$g_{k:T}(x_{k:T}) = h_k(x_k) + g_{k+1:T}(x_{k+1:T})$$

for some functions  $\{h_i\}_{i=1}^T$ . That is, we cannot decompose  $K$  recursively into an immediate complexity plus a future complexity. In particular, we cannot obtain  $K(x_{1:T})$  by backward induction.<sup>7</sup>

See [33] for proof. Due to Proposition 1, we cannot naively apply dynamic programming to solve (2) (for large enough  $T$ ). More precisely, the objective function

$$\Psi(a_{0:T}) := \sum_{t=0}^T r_t(s_t, a_t) - \beta K(a_{0:T}),$$

subject to the dynamics  $s_{t+1} = f_t(s_t, a_t)$  and start state  $s_0 \in S$ , cannot be decomposed on the form  $\Psi(a_{k:T}) = \Psi_0(a_0) + \Phi_{1:T}(a_{1:T})$ , where  $\{\Phi_{k:T}\}_{k=1}^T$  are given recursively by  $\Phi_{k:T}(x_{k:T}) = \Psi_k(x_k) + \Phi_{k+1:T}(x_{k+1:T})$  for some functions  $\{\Psi_k\}_{k=0}^T$ . Indeed, if such a decomposition existed, it would contradict Proposition 1. Thus,  $\Psi(a_{0:T})$  cannot be maximised using dynamic programming. Fortunately, there are methods that circumvent this issue, presented next.

## IV. COMPLEXITY-AWARE PLANNING ALGORITHMS

This section answers question 3 in the problem statement, presenting two algorithms that circumvent the dynamic programming issue given by Section III-C. The first algorithm, COPS, restricts the task to find simple optimal policies, while the second algorithm, SCAP, modifies the objective focusing on local (stage-wise) complexity.

### A. Complexity-guided Optimal Policy Search (COPS)

COPS seeks a simple optimal policy as in Example 1 by a two-step procedure. Step 1 conducts ordinary dynamic programming (maximising the total reward  $\sum_{t=0}^T r_t(s_t, a_t)$ ). This yields a mapping  $\Pi : \mathbb{T} \times S \rightarrow \mathcal{P}(A)$  such that  $\Pi(t, s) \subseteq A$  are the optimal actions at time  $t \in \mathbb{T}$  and state  $s \in S$ .<sup>8</sup> In step 2, a uniform-cost search [29] is executed to find an optimal action sequence  $a_{0:T}$  with low execution complexity. A node  $n$  in this search is on the form  $n = (t, s_t, a_{0:t-1}) \in \mathbb{T} \times S \times A^t$ , where  $t$  and  $s_t$  are the current time and state and  $a_{0:t-1}$  is the sequence of

<sup>7</sup>Note that  $h_i$  and  $g_{i:T}$  can be any functions, not necessarily computable. If one restricts these functions to be computable, then the result follows easily from the fact that  $K$  itself is not commutable.

<sup>8</sup>Here,  $\mathcal{P}(A)$  is the power set of  $A$ , i.e., the set of all subsets of  $A$ .

previous actions taken to arrive at state  $s_t$  at time  $t$ . If node  $n$  is not a terminal node (i.e.,  $t \leq T$ ), the children of  $n$  are given by all  $(t+1, s_{t+1}, a_{0:t})$  such that  $a_t \in \Pi(t, s_t)$  and  $s_{t+1} = f_t(s_t, a_t)$ , i.e., we expand only over optimal actions. The cost for a node  $n = (t, s_t, a_{0:t-1})$  is set to  $c(n) = K(a_{0:t-1})$ . The heuristic intuition behind the cost is that a low-complexity sequence  $a_{0:T}$  should be more likely to have low-complexity subsequences  $a_{0:t-1}$  and, thus, the cost focuses the search on low-complexity sequences, enabling moderate search depths.

The uniform-cost search is conducted by iteratively generating children of the node  $n$  with the lowest cost, starting from the root node  $n_0 = (0, s_0, \emptyset)$ . We terminate the search when a terminal node  $n = (T+1, s_{T+1}, a_{0:T})$  has the lowest cost and return its action sequence  $a_{0:T}$ . The algorithm is summarised by Algorithm 1. The termination is motivated by the following result:

*Proposition 2:* Assume  $K(a_{0:t-1}) \leq K(a_{0:t})$  holds for all  $t \leq T$  and optimal action sequences  $a_{0:T}$ . Then Algorithm 1 returns an optimal action sequence  $a_{0:T}$  with lowest execution complexity, i.e.,  $a_{0:T}$  maximises (3).

We stress that  $K(a_{0:t-1}) \leq K(a_{0:t})$  is not always true, since adding  $a_t$  to  $a_{0:t-1}$  may result in higher regularity than  $a_{0:t-1}$  has alone. However,  $K(a_{0:t-1}) \leq K(a_{0:t})$  is more anticipated since  $K$  is on average an increasing function with respect to sequence length [19], motivating the assumption in Proposition 2, and why the algorithm can work well in practice. Proposition 2 follows readily by applying uniform-cost search properties, see [33].

Efficient searching in Algorithm 1 is possible for moderate horizon lengths, demonstrated numerically in Section V-B. However, for longer horizons, the procedure becomes intractable. The next algorithm works for longer horizons by focusing on local (stage-wise) complexity instead of the complexity of the whole action sequence. COPS can be seen as a special case of this latter algorithm with only one stage and  $\beta_0 > 0$ , defined below in (4), sufficiently low.

---

#### Algorithm 1: COPS

---

**Input:** System  $\mathcal{M}$  as in (1) and start state  $s_0 \in S$ .

**Output:** Low-complexity optimal action sequence  $a_{0:T}$ .

**Step 1: Dynamic programming**

$V_{T+1}(s) = 0$  for all  $s \in S$ ;

**for**  $t = T, T-1, \dots, 0$  **do**

**forall**  $s \in S$  **do**

**forall**  $a \in A$  **do**

$Q_t(s, a) := r_t(s, a) + V_{t+1}(f_t(s, a));$

$V_t(s) = \max_{a \in A} Q_t(s, a);$

$\Pi(t, s) = \arg \max_{a \in A} Q_t(s, a);$

**Step 2: Uniform-cost search**

$q = \{\}$  # priority queue ordered by the cost  $c$ ;

Append root node  $n_0 = (0, s_0, \emptyset)$  to  $q$ ;

**loop do**

    Pop first node  $n = (t, s_t, a_{0:t-1})$  from  $q$ ;

**if**  $t=T+1$  **then**

**return**  $a_{0:t-1}$ ;

**forall**  $a_t \in \Pi(t, s_t)$  **do**

        Append node  $n = (t+1, f_t(s_t, a_t), a_{0:t})$  to  $q$ ;

---

## B. Stage-Complexity-Aware Program (SCAP)

SCAP modifies the objective to focus only on local (stage-wise) complexity. More precisely, the modified objective is set to

$$\max_{a_{0:T}} \left[ \sum_{t=0}^T r_t(s_t, a_t) - \sum_{k=0}^K \beta_k K(\mathbf{a}_k) \right]. \quad (4)$$

Here,  $K+1$  is the number of stages in the horizon partition, each with stage length  $l$ , and  $\mathbf{a}_k := a_{lk:l(k+l)-1} \in A^l$  is the executed action sequence at stage  $k \in \{0, \dots, K\}$ , penalised by its execution complexity  $K(\mathbf{a}_k)$  with weight  $\beta_k \geq 0$ .

The key insight is that the modified objective enables dynamic programming over the stages. More precisely, initialise the value function as  $V_{K+1}(s) = 0$  for all  $s \in S$  and obtain the value function for the remaining stages using the backward recursion

$$V_k(s_{lk}) = \max_{\mathbf{a}_k \in A^l} [r_k^l(s_{lk}, \mathbf{a}_k) - \beta_k K(\mathbf{a}_k) + V_{k+1}(f_k^l(s_{lk}, \mathbf{a}_k))]. \quad (5)$$

Here,  $f_k^l(s_{lk}, \mathbf{a}_k)$  denotes the state one arrives at by sequentially applying  $\mathbf{a}_k$  starting from  $s_{lk}$ , and  $r_k^l(s_{lk}, \mathbf{a}_k) = \sum_{t=lk}^{l(k+l)-1} r_t(s_t, a_t)$ . Provided  $l$  is small enough, the maximisation in (5) can be done by going through all  $\mathbf{a}_k \in A^l$ .

1) *Hard-constrained version:* An alternative to (4) is the hard-constraint objective

$$\max_{a_{0:T}} \sum_{t=0}^T r_t(s_t, a_t), \quad \text{s.t. } K(\mathbf{a}_k) \leq L_k, \quad \forall k \quad (6)$$

for some constants  $L_k \geq 0$ . In this case, we conduct dynamic programming with  $V_{K+1}(s) = 0$  for all  $s \in S$  and backward recursion

$$V_k(s_{lk}) = \max_{K(\mathbf{a}_k) \leq L_k} \left[ r_k^l(s_{lk}, \mathbf{a}_k) + V_{k+1}(f_k^l(s_{lk}, \mathbf{a}_k)) \right]. \quad (7)$$

Using (7) enables more efficient computations than (5) since the sequences  $\mathcal{A}_k := \{\mathbf{a}_k \in A^l : K(\mathbf{a}_k) \leq L_k\}$  are typically only a fraction of  $A^l$ . Moreover,  $\mathcal{A}_k$  can be computed beforehand by going through all sequences in  $A^l$  (tractable for small  $l$ ), or sought using a uniform-cost search (for moderate  $l$ , see [33] for details).

2) *Action sequence extraction:* Once  $V_k$  has been computed using (5) or (7),  $a_{0:T}$  maximising (4) or (6) can be readily obtained, for a given start state  $s_0 \in S$ , by forward simulation over the stages (see [33] for details).

Algorithm 2 summarises the procedure for the hard-constrained version.

## V. NUMERICAL EVALUATIONS

This section presents case studies evaluating the complexity-aware planning algorithms proposed in Section IV. We describe the test environment in Section V-A, evaluate COPS in Section V-B and then SCAP in Section V-C. The Kolmogorov complexity is estimated using the method from [40].

## Algorithm 2: SCAP

---

**Input:** System  $\mathcal{M}$  as in (1) and start state  $s_0 \in S$ .  
**Output:** Action sequence  $a_{0:T}$  maximising (6).  
**Step 1: Dynamic programming**  
Set  $V_{T+1}(s) = 0$  for all  $s \in S$ ;  
Compute  $\mathcal{A}_k = \{a_{0:l-1} \in A^l : K(a_{0:l-1}) \leq L_k\}$  for  $k = 0, 1, \dots, K$ ;  
**for**  $k = K, K-1, \dots, 0$  **do**  
  **forall**  $s \in S$  **do**  
    **forall**  $a_{0:l-1} \in \mathcal{A}_k$  **do**  
       $Q_k(s, a_{0:l-1}) = r_k^l(s, a_{0:l-1}) + V_{k+1}(f_k^l(s, a_{0:l-1}))$ ;  
       $V_k(s) = \max_{a_{0:l-1}} Q_k(s, a_{0:l-1})$ ;  
    **endforall**  
  **endforall**  
**Step 2: Action sequence extraction**  
**for**  $k = 0, 1, \dots, K$  **do**  
  Set  $a_{lk:l(k+l)-1} \in \arg \max_{a_{0:l-1} \in \mathcal{A}_k} Q_k(s_{lk}, a_{0:l-1})$ ;  
   $s_{l(k+1)} = f_k^l(s_{lk}, a_{lk:l(k+l)-1})$ ;  
  **return**  $a_{0:T}$

---

## A. Test Environment

As test environment, we consider a system  $\mathcal{M}$  as in (1) where a robot (the agent) moves around in a room with discrete steps. More precisely, the room is a square with coordinates  $N = \{1, \dots, n\}$  in each direction forming the state space  $S = N \times N$ . At each time, the robot can move to any of the neighbouring spots or stay, i.e., the actions are  $A = \{(\pm 1, 0), (0, \pm 1), (0, 0)\}$  with dynamics  $s_{t+1} = s_t + a_t$ . If the robot executes an action that would move it outside the room, then it stays at the same spot. The reward is zero everywhere except for all state-action pairs  $(s, a)$  that takes the robot to a given goal state  $s^*$ , i.e.,  $f(s, a) = s^*$ , with reward  $r(s, a) = 1$ . We set the goal state  $s^* = (n, n)$  to one of the corners of the room. The robot starts at  $s_0 = (1, 1)$  and the horizon  $T = 2(n-1) - 1$  is set so that the agent precisely receives a reward for reaching  $s^*$  if acting optimally.

## B. Evaluation 1

1) *Setup:* We first consider COPS. To see how output and calculation time changes with problem size, we study the test environment  $\mathcal{M}$  with different room sizes  $n = 10, 20$  and  $60$ , called the small, medium and large room, respectively. We run the search until it has found 30 low-complexity action sequences, by keep expanding the node tree. Concretely, this is done by modifying the if-statement in Algorithm 1 to a condition appending every found action sequence  $a_{0:T}$  into a list until this list has 30 sequences, and then return the list.

2) *Result:* The running time for the three rooms are 21 seconds, 16 minutes and 15 minutes, respectively, and corresponding trajectories to the found action sequences are given in Fig. 1a, 1b and 1c, with execution complexities in Tables Ia, Ib, and Ic.

As a first example, consider the trajectory of the first action sequence found in the small room, labeled 1 in Fig. 1a. Here, the robot goes right until it reaches the upper right corner and then goes down to  $s^*$ . That is, it exploits executing the same actions in batches to lower complexity, reaching an execution complexity of 47.30 seen in Table Ia. We also note that this is the same complexity as action sequences 2-4 in Fig. 1a.

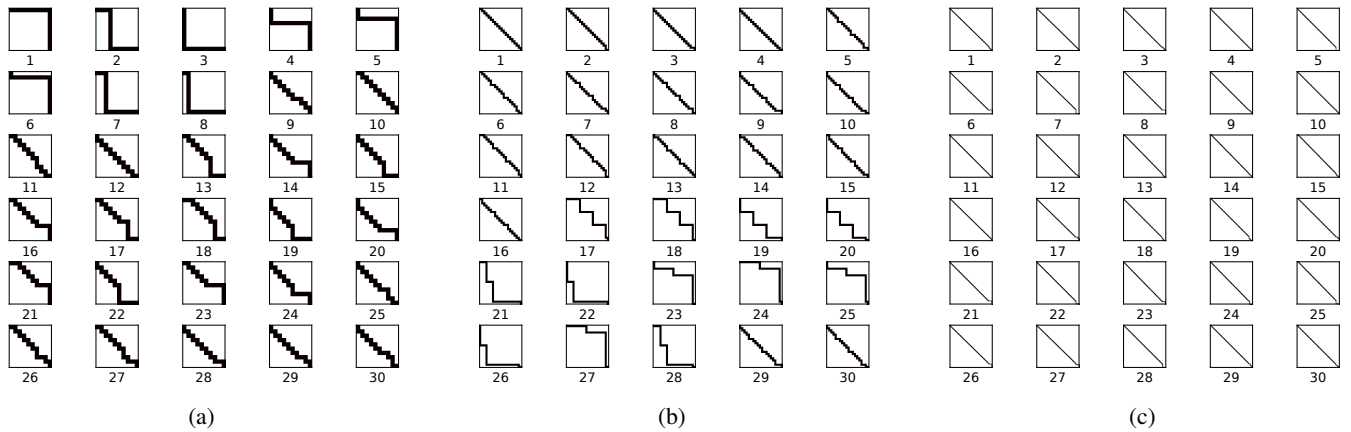


Fig. 1: Trajectories of the first 30 action sequences using COPS for: (a) small room, (b) medium room, and (c) large room.

TABLE I: Execution complexity for the found action sequences corresponding to Figure 1.

(a)		(b)		(c)	
Action sequence	Execution complexity	Action sequence	Execution complexity	Action sequence	Execution complexity
1-4	47.30	1-4	36.49	1-4	58.80
5-8	47.79	5-8	38.39	5-8	59.96
9-12	47.91	9-12	38.79	9-16	60.53
13-16	47.92	13-16	38.39	17-20	60.88
17-24	48.30	17-20	38.60	21-28	61.20
25-30	48.36	21-28	62.57	29-30	61.64
		29-30	39.75		

That 1 and 3 have the same complexity is due to symmetry, and the same is true for 2 and 4. However, why e.g., 1 and 2 have the same complexity (apart from the intuition that they both look like low-complexity executions) is unclear. It could be an inherited feature from the Kolmogorov complexity, or a bias from the estimation method, see [33] for a discussion of the latter.

Looking now at all sequences in all three cases, we see that it is in general common to find sequences executing batches of the same actions and then alternate between such batches to lower execution complexity (1-8 in Fig. 1a, and 17-28 in Fig. 1b). Another typical feature to lower complexity is to exploit 2-periodic alternation between going down and going right (9-30 in Fig. 1a, 1-16 and 29-30 in Fig. 1b, and 1-30 in Fig. 1c). Also, what type of execution with lowest complexity varies with the room size, i.e., is horizon-dependent. Once again, this could come from the Kolmogorov complexity itself, where varying sequence length may facilitate different compression techniques, but could also originate from the estimation method.

We also note in Tables Ia to Ic that the algorithm mainly finds sequences in increasing complexity-order as anticipated by Proposition 2. Indeed, only the medium room case causes some unordered sequences, where notably the search finds the higher complexity sequences 21-28 with execution complexity 62.57 before reaching sequences 29-30 with complexity 39.75. This detour, caused by a violation in the assumption in Proposition 2, explains the long search time the medium room yields, longer than the large room although the latter has a larger horizon.

Finally, the cost  $c(n) = K(a_{0:t-1})$  enables moderate horizon lengths  $T$ . In particular, for the large room, the number of possible optimal action sequences is around  $2^{(T+1)} \approx 10^{36}$  after dynamic programming. However, the search only iterates  $10^6$  nodes to find the first low-complexity sequence, a significant decrease due to the complexity-guiding cost.

### C. Evaluation 2

We continue with SCAP, considering the hard-constrained version given by Algorithm 2, calculating  $\mathcal{A}_k$  by going through all sequences. The aim of the investigation is mainly to see how the complexity-performance tradeoff affects the behaviour. Towards this, we consider two similar setups, both considering the large room but where the first setup has the goal state  $s^*$  in the lower right corner of the room and the second setup has  $s^*$  in the middle of the room.

1) *Setup 1:  $s^*$  in the corner:* In Setup 1, we penalise each stage equally with complexity limit  $L_k = L$ , having stage length  $l = 12$ , and consider the large room from Section V-B with a slightly larger horizon  $T = 119$  to be compatible with the stage partition (i.e.,  $l \cdot (K + 1) = T + 1$ ), setting  $K = 9$  accordingly. The change in tradeoff comes into play by varying  $L$  with values 26, 28 and 30. These values are picked since they illustrate the tradeoff well, allowing almost no action sequences for  $L = 26$  while many for  $L = 30$ .

2) *Result:* Value functions  $V_0(s)$  are plotted in Fig. 2 for  $L \in \{26, 28, 30, \infty\}$ , where  $L = \infty$  is as a reference obtained by ordinary dynamic programming maximising the total reward  $\sum_{t=0}^T r_t(s_t, a_t)$ . Blue (red) colour encodes low (high) value.




Fig. 2:  $V_0(s)$  for Setup 1 in Section V-C and  $L \in \{26, 28, 30, \infty\}$ . Shown are also two corresponding trajectories (for each finite  $L$ ) starting from A and B, obtained by Algorithm 2.

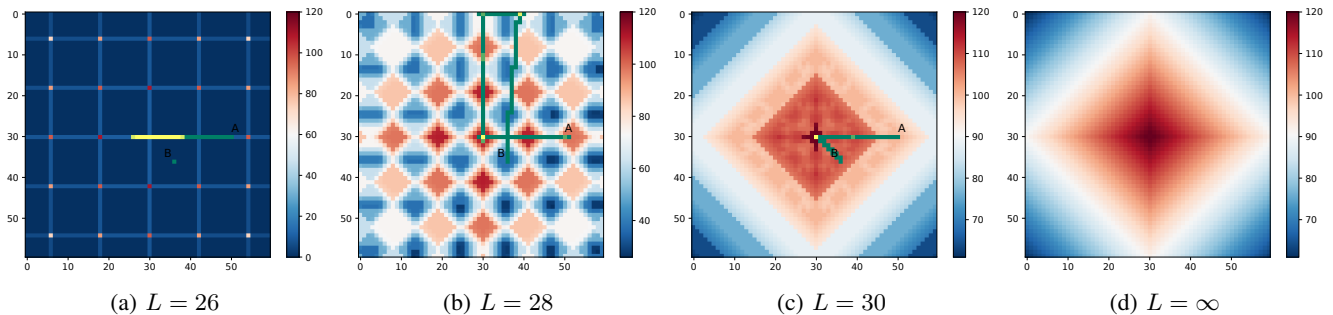


Fig. 3: Same as Fig 2. except for Setup 2. Shown are also two trajectories starting from A and B, obtained by Algorithm 2.

For  $L = 26$ , the room is divided into squares of length  $l = 12$  due to the stage partition and the low complexity limit  $L = 26$ . More precisely, at each stage  $k$ ,  $\mathcal{A}_k$  consists of only five action sequences: the ones executing only one action throughout the stage. This constraint implies that the robot sometimes has to wait at the walls, yielding the discontinuous jumps in the value function seen in Fig. 2 (a). Consider for example the two trajectories in Fig. 2 (a) generated by Algorithm 2, coloured in a green-to-yellow scale, where states more visited are more yellow. The robot starting at A reaches  $s^*$  without waiting. The robot starting at B goes right two stages and then down two stages, but hits the wall at its second going-right stage before the stage is done. This causes waiting at the wall and thus longer time to reach  $s^*$ , even though B is actually closer to  $s^*$  than A.

For  $L = 28$ , the squares in Fig. 2 (b) are greatly diminished compared to  $L = 26$ , due to additional possible action sequences in  $\mathcal{A}_k$ . Here, a robot starting at B does not wait anymore at the wall. We note also that the squares are more evident near the goal state since the robot has less time here to adapt and avoid wall waiting.

Finally, for  $L = 30$ ,  $\mathcal{A}_k$  has increased so much that the difference with the unconstrained case is minor, and the trajectories from A and B, already being optimal at  $L = 28$ , has not changed.

3) *Setup 2:  $s^*$  in the middle:* Setup 2 is identical to Setup 1, except that  $s^*$  is placed in the middle of the room. Thus, the robot can no longer heavily rely on the wall dynamics as in Setup 1. This notably changes the outcome.

4) *Result:* Similar to Setup 1, we plot  $V_0(s)$  for different  $L$  in Fig. 3.

For  $L = 26$ ,  $V_0(s)$  has higher values in a grid-like pattern, with highest value in the middle of the room at  $s^*$  and high values at points being away a multiple of  $l = 12$  from  $s^*$ . At all these points, the robot can (using the five sequences in  $\mathcal{A}_k$ ) reach and stay at  $s^*$  resulting in a high total reward. Between those points are also states with higher value. Here, the robot can reach but not stay at  $s^*$ ; Instead, the robot oscillates back and forth crossing the  $s^*$  multiple times. This is seen for the robot starting at A in Fig. 3 (a). For the remaining states, the robot can never reach the goal state, due to the heavy complexity constraint. This is the case for the robot starting and staying at B. Thus, for low complexity limits, the robot may never reach the intended objective.

As we increase  $L$  to 28, the grid-like pattern is expanded to groups of states due to the increased set of admissible action sequences. A robot starting at A reaches now and stays at  $s^*$  as can be seen in Fig. 3 (b). The robot starting at B also reaches and stays at  $s^*$ , but takes a detour to be within the local complexity limit, reaching the upper wall and then down to  $s^*$ . Thus, for low complexity limits, the robot may execute an action sequence which is locally of low complexity, but globally quite complex.

Finally, for  $L = 30$ ,  $V_0(s)$  is similar to the unconstrained case, except that the stage length  $l = 12$  together with the complexity constraint still partitions the space into entities of length  $l = 12$ , similar to Fig. 2 (a) (but in this case into diamond-like shapes). The robots starting at A and B now reaches the goal fast with the high complexity limit.

## VI. CONCLUSION

In this paper, we have considered complexity-aware planning for DFA with rewards as outputs. We first defined a complexity measure for deterministic policies based on Kolmogorov complexity. Kolmogorov complexity is used since it can detect computational regularities, a typical feature for optimal policies.

We then introduced a complexity-aware planning objective based on our complexity measure, yielding an explicit trade-off between a policy's performance and complexity. It was proven that maximising this objective is non-trivial in the sense that dynamic programming is infeasible.

We presented two algorithms for obtaining low-complexity policies, COPS and SCAP. COPS finds a policy with low complexity among all optimal policies, following a two-step procedure. In the first step, the algorithm finds all optimal policies, without any complexity constraints. Dropping the complexity constraints, this step can be done using dynamic programming. The second step runs a uniform-cost search over all optimal policies, guided by a complexity-driven cost, favouring low-complexity policies. SCAP modifies instead the objective penalising policies locally for executing complex manoeuvres. This is done by partitioning the horizon into stages and enforce local complexity constraints over these stages, where the partition enables dynamic programming. We illustrated and evaluated our algorithms on a simple navigation task where a mobile robot tries to reach a certain goal state. Our algorithms yield low-complexity policies that concur with intuition.

Future work includes comparisons with other estimation methods of the Kolmogorov complexity and how uncertainty (e.g., stochasticity) and feedback (e.g., receding horizon) can be incorporated into the existing framework. Finally, another challenge is to tractably maximise the complexity-aware planning objective in the general case.

## REFERENCES

- [1] D. Abel, et al. State abstraction as compression in apprenticeship learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3134–3142, 2019.
- [2] J. Aslanides, et al. Universal reinforcement learning algorithms: Survey and experiments. *arXiv preprint arXiv:1705.10557*, 2017.
- [3] F. Attneave. Some informational aspects of visual perception. *Psychological review*, 61(3):183, 1954.
- [4] H. B. Barlow. Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1:217–234, 1961.
- [5] R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [6] O. Biza, et al. Learning discrete state abstractions with deep variational inference. *arXiv preprint arXiv:2003.04300*, 2020.
- [7] M. Booker and A. Majumdar. Learning to actively reduce memory requirements for robot control tasks. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, pages 125–137, 2021.
- [8] N. Chmait, et al. Factors of collective intelligence: How smart are agent collectives? In *Proceedings of the Twenty-second European Conference on Artificial Intelligence*, pages 542–550, 2016.
- [9] M. K. Cohen, et al. A strongly asymptotically optimal agent in general environments. *arXiv preprint arXiv:1903.01021*, 2019.
- [10] F. Delmotte, T. R. Mehta, and M. Egerstedt. A software tool for hybrid control. *IEEE Robotics Automation Magazine*, 15(1):87–95, 2008.
- [11] M. B. Egerstedt and R. W. Brockett. Feedback can reduce the specification complexity of motor programs. *IEEE Transactions on Automatic Control*, 48(2):213–223, 2003.
- [12] R. Fox and N. Tishby. Minimum-information LQG control part i: Memoryless controllers. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 5610–5616. IEEE, 2016.
- [13] A. Girard and G. J. Pappas. Approximate bisimulation: A bridge between computer science and control theory. *European Journal of Control*, 17(5-6):568–578, 2011.
- [14] J. E. Hopcroft, et al. *Introduction to Automata Theory, Languages, and Computation (3rd Edition)*. Addison-Wesley, 2006.
- [15] M. Hutter. *Universal artificial intelligence: Sequential decisions based on algorithmic probability*. Springer Science & Business Media, 2004.
- [16] M. J. Kearns, Y. Mansour, and A. Y. Ng. Approximate planning in large POMDPs via reusable trajectories. In *Advances in Neural Information Processing Systems*, pages 1001–1007, 2000.
- [17] D. T. Larsson, D. Maity, and P. Tsiontras. An information-theoretic approach for path planning in agents with computational constraints. *arXiv preprint arXiv:2005.09611*, 2020.
- [18] A. Lempel and J. Ziv. On the complexity of finite sequences. *IEEE Transactions on information theory*, 22(1):75–81, 1976.
- [19] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, 4th edition, 2020.
- [20] L. Liu, A. Chattopadhyay, and U. Mitra. On exploiting spectral properties for solving MDP with large state space. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1213–1219. IEEE, 2017.
- [21] George H. Mealy. A method for synthesizing sequential circuits. *The Bell System Technical Journal*, 34(5):1045–1079, 1955.
- [22] V. Mnih, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [23] J. M. O’Kane and D. A. Shell. Concise planning and filtering: Hardness and algorithms. *IEEE Transactions on Automation Science and Engineering*, 14(4):1666–1681, 2017.
- [24] A. Pervan and T. D. Murphey. Algorithmic design for embodied intelligence in synthetic cells. *IEEE Transactions on Automation Science and Engineering*, 18(3):864–875, 2021.
- [25] B. Recht. A tour of reinforcement learning: The view from continuous control. *arXiv preprint arXiv:1806.09460*, 2018.
- [26] N. Roy and G. J. Gordon. Exponential family PCA for belief compression in POMDPs. In *Advances in Neural Information Processing Systems*, pages 1667–1674, 2003.
- [27] J. Rubin, O. Shamir, and N. Tishby. Trading value and information in MDPs. In *Decision Making with Imperfect Decision Makers*, pages 57–74. Springer, 2012.
- [28] S. Russell. Rationality and intelligence: A brief update. In *Fundamental issues of artificial intelligence*, pages 7–28. Springer, 2016.
- [29] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall Press, USA, 3rd edition, 2009.
- [30] A. A. Rusu, et al. Policy distillation. *arXiv preprint arXiv:1511.06295*, 2015.
- [31] D. Silver, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [32] C. R. Sims. Efficient coding explains the universal law of generalization in human perception. *Science*, 360(6389):652–656, 2018.
- [33] E. Stefansson and K. H. Johansson. Computing complexity-aware plans using Kolmogorov complexity. *arXiv preprint arXiv:2109.10303*, 2021.
- [34] T. Tanaka, P. M. Esfahani, and S. K. Mitter. LQG control with minimum directed information: Semidefinite programming approach. *IEEE Transactions on Automatic Control*, 63(1):37–52, 2017.
- [35] T. Tanaka, et al. Transfer-entropy-regularized Markov decision processes. *arXiv preprint arXiv:1708.09096*, 2017.
- [36] N. Tishby and D. Polani. Information theory of decisions and actions. In *Perception-action cycle*, pages 601–636. Springer, 2011.
- [37] S. Udrescu, et al. AI Feynman 2.0: Pareto-optimal symbolic regression exploiting graph modularity. *arXiv preprint arXiv:2006.10782*, 2020.
- [38] H. S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):131–147, 1968.
- [39] Y. Ho and K. Chu. Team decision theory and information structures in optimal control problems—part i. *IEEE Transactions on Automatic Control*, 17(1):15–22, 1972.
- [40] H. Zenil, et al. A decomposition method for global evaluation of Shannon entropy and local estimations of algorithmic complexity. *Entropy*, 20(8):605, 2018.
- [41] K. J. Åström. *Introduction to Stochastic Control Theory*, volume 70 of *Mathematics in science and engineering*. Academic Press, 1970.