TERC: A Transfer Entropy Redundancy Criterion for State Variable Selection in Reinforcement Learning

Anonymous authors
Paper under double-blind review

Abstract

Identifying the most suitable variables to represent the state is a fundamental challenge in Reinforcement Learning (RL). These variables must efficiently capture the information necessary for making optimal decisions. In order to address this problem, in this paper, we introduce the Transfer Entropy Redundancy Criterion (TERC), an information-theoretic criterion, which determines if there is *entropy transferred* from observable state variables to actions during training. We define an algorithm based on TERC that provably excludes variables from the observable state that do not affect the agent's policy during learning. This yields compact state representations that reduce inference time by up to 2.6×. Our approach is policy-dependent, making it agnostic to the underlying learning algorithm. The efficiency gains we demonstrate arise at retraining and inference time on the reduced state.

Our method improves both retraining and inference efficiency. We demonstrate its effectiveness across three distinct algorithm classes, namely tabular Q-learning, Actor-Critic, and Proximal Policy Optimization (PPO), evaluated in a range of environments. Furthermore, to highlight the differences between the proposed methodology and the current state-of-the-art feature selection approaches, we present a series of controlled experiments on synthetic data, before generalizing to real-world decision-making tasks. We also introduce a representation of the problem that compactly captures the transfer of information from observable state variables to actions as Bayesian networks.

1 Introduction

The choice of an appropriate state representation remains a key design challenge of every Reinforcement Learning (RL) system. This process involves finding simplified representations of the information required to learn optimal policies. In order to achieve this, various techniques have been proposed (Finn et al., 2016; Laskin et al., 2020; Gelada et al., 2019; Stooke et al., 2021). Nevertheless, several issues can be identified with these approaches. Firstly, most of these methods fail to provide a mechanism by which the state's dimensionality can be reduced. Notable exceptions include the state-variable selection of Hein et al. (2018) and the knockoff approach of Ma et al. (2023). Consequently, at the time of deployment, most RL systems process unnecessarily large amounts of data. Second, they are used in a black-box fashion, providing no insight into the final form of the simplified representation. Finally, they fail to provide a solution for determining the appropriate state history lengths when temporally extended states are required, as demonstrated in Mnih et al. (2015); Pan et al. (2017). As a result of these shortcomings, most RL state design applications fall into the realm of intuition-guided heuristics (Dean & Givan, 1997; Ortiz et al., 2018; Reda et al., 2020; Liu et al., 2020). The use of such *ad hoc* methods results in state representations that incorporate non-informative variables, which might increase the training duration due to Bellman's curse of dimensionality (Bellman & Kalaba, 1959). They may also lead to state representations that lack the required information, preventing the agent from learning optimal policies.

In order to provide an intuition of the problem, let us consider the classic cart pole environment[1]. Imagine a cart that can move left or right along a track. On top of this cart, there is a pole standing upright. The goal is to keep the pole balanced and prevent it from falling over. You can do this by moving the cart left or right to keep the pole centered. The challenge is to figure out the best way to move the cart to keep the pole balanced for as long as possible. Typically, at time t, this environment is characterized by a 4-dimensional set of variables $s^t = [x^t, \theta^t_{pole}, \dot{x}^t, \dot{\theta}^t_{pole}]$, where x indicates the cart's x-position, θ^t_{pole} describes the angle of the pole, and $\dot{()}$ denotes the derivative with respect to time. Let us suppose that an additional random variable v_{rand}, which is independent of the physical system under consideration, is added to the environment. The resulting state will be $s^t_{extended} = [x^t, \theta^t_{pole}, \dot{x}^t, \dot{\theta}^t_{pole}, v_{rand}]$. v_{rand} is not informative for the problem at hand and a method to define optimal representations should be able to exclude it. Failure to remove such variables will lead to unnecessary computations once deployed[2]. Moreover, uninformative variables can impede the learning process of iterative decision-making systems (Grooten et al., 2023).

To address these issues, we propose the Transfer Entropy Redundancy Criterion (TERC) for the selection of the minimal set of the state variables. TERC is a criterion that allows us to determine *whether observable* state variables transfer entropy to actions during training (Schreiber, 2000). More specifically, TERC is based on the quantification of the reduction in uncertainty of the realizations of the policy when considering the set of observable state variables with and without the variable under consideration. This makes TERC a policy-dependent method, while being agnostic to the underlying learning algorithm. If this value is bigger than zero, the actions are said to depend on this state variable, and TERC is verified. Under these circumstances, the state variable is considered informative and included in our representation. Once all the informative observable state variables have been identified, our agent can be re-trained on the lightweight state, leading to greater efficiency at the time of deployment. We will provide the reader with a formal definition of TERC in Section 6.2[3].

Before describing the method in detail, we explicitly state the twofold objective of TERC: (i) *identifying a minimal sufficient state representation*, by providing a principled, information-theoretic criterion for which observable variables are necessary for the agent's learned policy; and (ii) *improving retraining and inference efficiency*, by retraining or deploying the policy on the resulting lower-dimensional state, which reduces inference cost and can improve sample efficiency during retraining.

In practical terms, we will introduce a set of methods that -in conjunction with the satisfaction of an assumption- will allow us to derive the minimal and optimal set of variables for state representation in presence of redundant and synergistic relationships. Even when this assumption does not hold, our method still identifies an optimal, though not necessarily minimal, representation. This set is constructed starting from an empty set and adding variables to it by verifying TERC. Indeed, the core principle of our methodology is to include only the variables in the state representation that are informative. We will show that TERC can achieve this even in the presence of perfectly redundant state variables (Frye et al., 2020; Kumar et al., 2020). Consequently, we will show that it is possible to obtain an optimal subset whose realizations reduce the entropy of the actions identically to the original set of all observable state variables. Additionally, we will demonstrate that this method can serve a dual purpose: not only does it help in omitting uninformative variables from our representation, but it also facilitates investigating how variable importance varies as the agent trains, thereby enhancing the overall interpretability of the learning process.

In addition to discussing the theoretical basis of TERC, to further validate our approach from a practical point of view, we will present an extensive experimental evaluation through a selection of novel and existing environments. We will firstly apply TERC to purely synthetic data, considering different complex relationships between the observable variables. We will then introduce the 'Secret Key Game', a novel environment in which an agent is tasked with learning a secret message from an observable state comprising not only 'secret keys' (from which the secret message is calculated), but also superfluous 'decoy keys'. The calculation

[1]Many implementations of this environment exist, like that implemented in the now publicly maintained Gym/Gymnasium framework (`https://gymnasium.farama.org/environments/classic_control/cart_pole/`). In the evaluation of TERC presented in a later section, we use the original implementation by OpenAI (Brockman et al., 2016).

[2]Indeed, this is of critical importance given the increasing energy footprint of today's machine learning systems (Patterson et al., 2022).

[3]The code for the derivation of the measures at the basis of TERC is available at this URL: [*We plan to release the TERC library and the code used for all the simulations presented in this manuscript upon its acceptance*].

of the secret message is based on a classic method for secure multiparty communication, as described in (Shamir, 1979; Blakley, 1979). We will then consider a set of physics-based examples. We will consider three very popular Gym environments: Cart Pole, Lunar Lander, and Pendulum. Finally, we will illustrate the versatility of TERC, applying it to strategic sequential games, in which the temporal dimension plays a key role. In particular, we will focus on the problem of learning optimal state history lengths when playing against a Tit-For-N-Tats (TFNT) opponent in the Iterated Prisoner's Dilemma.

Summary of contributions. The contributions of this paper can be summarized as follows:

- We propose TERC, an effective criterion for the selection of state variables in reinforcement learning, discussing its theoretical foundations in detail.
- We analyze potential issues arising from the presence of perfect conditionally redundant variables in the observable state and we introduce algorithmic solutions to deal with them.
- We corroborate our theoretical results experimentally by means of an extensive evaluation considering a variety of RL environments.
- Finally, we discuss the extent to which TERC enhances the interpretability of RL systems and its practical applications.

2 Related Work

As discussed in the introduction, the method presented in this paper is related to the field of feature selection in machine learning. The proposed solution is then applied to the problem of state representation in RL. Consequently, we will first review the state of the art in the area of feature selection methods before discussing the relevant work on state representation learning.

Feature selection. Our objective is to devise a measure that assesses how informative state variables are given an agent's actions. This is conceptually similar to the idea of 'feature importance' in the field of feature selection (Ribeiro et al., 2016; Shrikumar et al., 2017; Plumb et al., 2018; Sundararajan & Najmi, 2020; Chen et al., 2020). The exact definition of 'feature importance' is an open problem *per se* (Catav et al., 2021; Janssen et al., 2023). However, it is generally agreed that its practical objective is to rank the input variables of a machine learning model based on their ability to predict the model output. For example, Shapley values, originally introduced in a game-theoretic context for allocating resources in cooperative systems (Shapley, 1953), have been repurposed as indicators of feature importance. One of the first examples was Multi-perturbation Shapley value Analysis (MSA) (Keinan et al., 2004; Cohen et al., 2007), which entails considering features as agents which cooperate to predict model outcomes. The Shapley Additive Explanation (Lundberg & Lee, 2017) and Shapely Additive Global Explanation (Lundberg & Lee, 2017) models expanded on this, allowing the ranking of input features both globally (predicting all model outputs) and locally (predicting a single model output). Despite the success of Shapley value-based methods (Apley & Zhu, 2020; Covert et al., 2020; Kwon & Zou, 2022), both theoretical and experimental findings have shown that the presence of redundant information can invalidate these methods' results (Frye et al., 2020; Kumar et al., 2020). To deal with these issues, in this work we draw inspiration from information-theoretic filter feature selection methods (Battiti, 1994; Tesmer & Estévez, 2004; Peng et al., 2005; Brown et al., 2012; Gao et al., 2016; Chen et al., 2018). Rather than ranking variables by feature importance, we propose the use of a 'target set' that we populate with variables that lead to desirable properties. This is analogous to what is proposed by Peng et al. (2005); Gao et al. (2016); Chen et al. (2018), except here we favorably derive the subset size, rather than requiring it as an arbitrary hyperparameter. Brown et al. (2012) proposed a method that did not suffer from such drawbacks, which has inspired subsequent works (Covert et al., 2023; Wollstadt et al., 2023). However, they fail to highlight optimal feature sets in some cases. Their methods add new features to the selected subset if they enhance the overall correlation with the target. This approach overlooks features that only become informative when considered in combination. Furthermore, such methods add features that maximize the correlation function. This function must be computed for each feature at every step, meaning the complexity of these techniques fails to scale linearly in time with

respect to the number of features (Borboudakis & Tsamardinos, 2019; Tsamardinos et al., 2019; Hao et al., 2021). Our method has the ability to evaluate the effects of complex relationships, such as redundancies and synergistic combinations, while remaining linear in time with respect to the number of features. In the specific RL setting, Hein et al. (2018) apply AMIFS-based feature selection (Tesmer & Estévez, 2004) to state variables; TERC differs by using transfer entropy to capture temporal/causal dependencies between states and actions, by handling conditional redundancy and synergy beyond pairwise interactions, and by providing theoretical guarantees of minimality. We also note the related line of work on RL-specific state-variable selection based on Sequential Knockoffs (SEEK) (Ma et al., 2023), which uses knockoff features to test the relevance of individual state variables.

Unsupervised derivation of state representations in reinforcement learning. The Infomax principle has been used as the basis for several unsupervised representation learning techniques; it involves the maximization of the mutual information (MI) between network inputs and appropriately constrained outputs (Linsker, 1988; Bell & Sejnowski, 1995). The methods relying on the Infomax principle have been successfully used to learn representations of natural language (Devlin et al., 2019), videos (Sun et al., 2019), images (Hjelm et al., 2019; Bachman et al., 2019), and RL states (Finn et al., 2016; Laskin et al., 2020). For high-dimensional and continuous input/output spaces, the associated MI is not computable, and, for this reason, the application of the Infomax principle is not practically possible (Song & Ermon, 2020). To address this issue, a variety of methods for estimating MI have been developed (van den Oord et al., 2019; Belghazi et al., 2018; Poole et al., 2019). The problem of defining a compact state representation has been studied primarily as an abstraction problem (James C. Bean, 1987; Lesort et al., 2018). The goal of state abstraction is to reduce the size of the associated state space by unifying the representation of the areas of the space that provide indistinguishable information. Examples of such techniques include bisimulation (Givan et al., 2003), homomorphism (Ravindran & Barto, 2003), utile distinction (McCallum & Ballard, 1996), and policy irrelevance (Jong & Stone, 2005). A complementary line of work uses recurrent neural networks to compress histories of observables into compact quasi-Markovian state representations for partially observable MDPs (Schäfer & Udluft, 2005; Schäfer et al., 2007). TERC is compatible with this perspective but operates at the level of variable selection rather than latent-state compression, and additionally derives the optimal state history length explicitly rather than learning it implicitly, as is done in Duell et al. (2010). These methods most commonly leverage MI-based abstraction principles (such as Infomax) to redefine the distributions associated with each of the original state variables, leading to new representations that preserve MI between sequential states, actions, or combinations thereof (Schwarzer et al., 2020); while concurrently disregarding irrelevant and redundant information. In our study, we take a different approach, focusing on the conditional relationships between state variables and actions during the learning phase. This allows us to eliminate variables that are not informative with respect to agent training. Upon completion of this process, we are not required to consider the entire set of variables, in contrast to the methods previously discussed. In other words, our solution derives a compact state representation with fewer variables. Once deployed, this reduced subset requires less computation for measuring, processing, and storing. Additionally, in the case of applications that require a temporally-extended state representation (e.g., a sequence of frames for an arcade game), our method is able to derive the optimal state history length.

3 Background and Notation

In the following, calligraphic symbols (e.g., $\mathcal{X}$) indicate sets, capital letters (e.g., X) represent random variables, and their realizations are denoted with lower-case letters (e.g., x). For a set $\mathcal{X}$ and subset $\mathcal{P} \subseteq \mathcal{X}$, we write $\mathcal{X}_{\setminus \mathcal{P}}$ to denote the set difference $\mathcal{X} \setminus \mathcal{P}$, and $\mathcal{X}_{\setminus (\mathcal{P}, \mathcal{P}')}$ as shorthand for $\mathcal{X} \setminus (\mathcal{P} \cup \mathcal{P}')$. We use Shannon's entropy, denoted $H(\cdot)$ (Shannon, 1948), to quantify uncertainty. We write $\mathscr{P}(\mathcal{X})$ for the power set of $\mathcal{X}$. A glossary of terminology and a summary of notation are provided in Appendix A.

Reinforcement learning. A Markov Decision Process (MDP) is characterized by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, R, T)$, consisting of the state space $\mathcal{S}$, action space $\mathcal{A}$, reward function R, and transition function T. An agent interacts with its environment by taking actions $a^t \in \mathcal{A}$ dependent on its current state $s^t \in \mathcal{S}$, generating trajectories $(s^1, a^1, s^2, a^2, \ldots, s^T, a^T)$, where each state is a vector $s^t = [x_1^t, x_2^t, \ldots, x_N^t]$. The agent learns a policy $\pi(a^t|s^t)$ that maximizes cumulative reward $J(\pi) = \mathbb{E}_\pi[\sum_{t=1}^{\infty} \gamma^t r^t]$, where $\gamma \in [0, 1]$

is a discount factor (Sutton & Barto, 2018). After training, we sample state-action pairs from trajectories to derive the set of random variables $\mathcal{X} = \{X_1, X_2, \ldots, X_N\}$ and action variable A, with distributions $p(\mathcal{X} = s) = \frac{1}{T}\sum_{t=1}^{T} p(\mathcal{X}^t = s)$ and $p(A = a) = \frac{1}{T}\sum_{t=1}^{T} p(A^t = a)$.

Information-theoretic concepts. Transfer entropy (TE) quantifies how knowledge of one variable reduces uncertainty about another, analogous to a nonlinear formulation of Granger causality (Granger, 1969; Schreiber, 2000). For variables X and Y, TE is defined as $TE_{Y\to X} = H(X^t|X^{t-1}) - H(X^t|X^{t-1}, Y^{t-1})$. When state or action spaces are high-dimensional or continuous, we estimate TE via mutual information (MI) using neural estimators (Belghazi et al., 2018). Two additional multivariate concepts are central to our analysis: *redundancy*, which arises when multiple variables provide overlapping information about a target, and *synergy*, which occurs when variables combine to provide information beyond their individual contributions (Williams & Beer, 2010). Both phenomena require careful treatment in feature selection, as standard correlation-based methods may fail to identify redundant variables or miss synergistic relationships entirely (Frye et al., 2020; Kumar et al., 2020).

4 Problem Statement and Training&Data-Collection Pipeline

Training and data-collection pipeline. The TERC workflow proceeds in three stages:

1. *Initial training.* Train an agent to convergence on the full observable state $s^t = [x_1^t, \ldots, x_N^t]$ (where $s^t \in \mathcal{S}$ and x_i is a realization of $X_i \in \mathcal{X}$) using any standard RL algorithm, yielding a policy π.
2. *Trajectory sampling.* Roll out π to collect state–action trajectories, and from them construct the empirical distributions of A and $\mathcal{X}$ (see Section 3).
3. *Variable selection.* Apply TERC to identify state variables on which the agent's actions exhibit zero dependence; these are excluded to obtain $\mathcal{X}_* \subseteq \mathcal{X}$, on which a new agent can be retrained for deployment.

Problem statement. We now introduce the problem statement in a formal way. Given the set of all observable state variables, $\mathcal{X} = \{X_1, X_2 \ldots X_N\}$, we aim to obtain the smallest possible subset $\mathcal{X}_* \subseteq \mathcal{X}$ from which RL agents can learn optimal policies.

Because TERC operates on distributions induced by π, the selected subset is defined *with respect to* π and may therefore depend on both the training algorithm and the resulting policy. Formally, this set is defined as follows[4]:

$$\mathcal{X}_* \in \{\mathcal{P} \in \mathscr{P}(\mathcal{X}) : (|\mathcal{P}| = \min_{H(A|\mathcal{P})=H(A|\mathcal{X})} |\mathcal{P}|) \quad \& \quad (H(A|\mathcal{P}) = H(A|\mathcal{X}))\}. \tag{1}$$

Since TERC only removes variables whose transfer entropy to actions is zero, the reduced state satisfies $H(A \mid \mathcal{X}_*) = H(A \mid \mathcal{X})$, so the reduced decision process retains the Markov property with respect to π and the convergence guarantees of Theorem 4(1) in Li et al. (2006). Additionally, diminishing the dimensionality of the state space results in a proportional reduction of the regret bounds (Yang & Wang, 2020). Our aim is to develop an approach to derive $\mathcal{X}_*$, which can then be used as the new state representation for agent training, such that $s_*^t = [x_1^t, x_2^t \ldots x_M^t]$, where $s_*^t \in \mathcal{S}_*$ and x_i^t is the realization of $X_i \in \mathcal{X}_*$.

5 Constrained Perfect Multivariate Conditional Redundancy

Constrained Perfect Multivariate Conditional Redundancy (CPMCR) is a condition that, if true, signifies the existence of two disjoint subsets $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{X}$ (where $\mathcal{P} \neq \mathcal{P}'$) that convey identical information about a target variable. When CPMCR holds, many existing feature importance methods fail to produce correct results (Breiman, 2001; Debeer & Strobl, 2020; Frye et al., 2020; Kumar et al., 2020). Two concrete cases illustrate this: if two observable state variables are identical (say $X_1 \equiv X_2$), then $\mathcal{P} = \{X_1\}$ and $\mathcal{P}' = \{X_2\}$ form trivially redundant subsets that satisfy CPMCR. Second, if X_1 and X_2 are independent binary variables

[4]It is worth noting that this set is defined as a member of a set of subsets, due to the existence of multiple representations that satisfy the conditions introduced in the problem statement.

with $X_3 = X_1 \oplus X_2$, then the single variable X_3 and the pair $\{X_1, X_2\}$ convey identical information about the XOR target. Our goal is to detect such redundant subsets and include only the smallest in our final state representation.

To formally define CPMCR, we first require that each element of a subset contributes meaningfully to its information content. For subset $\mathcal{P}'$, we define:

$$\psi_{\mathcal{P}'} = \Big(\forall P' \in \mathcal{P}' : H(Y|\mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}') < H(Y|\mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}'_{\setminus P'})\Big), \tag{2}$$

where $\mathcal{P}'_{\setminus P'}$ denotes subset $\mathcal{P}'$ with element P' removed. This condition ensures that every element of $\mathcal{P}'$ contributes to reducing uncertainty about Y; removing any single element increases conditional entropy. An analogous condition $\psi_{\mathcal{P}}$ applies to subset $\mathcal{P}$.

We now define CPMCR, denoted $\Psi(Y|\mathcal{X})$, as follows:

$$\begin{aligned}\Psi(Y|\mathcal{X}) = \Big(\exists \mathcal{P} \in \mathscr{P}(\mathcal{X}), \exists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus\mathcal{P}}) :& \\ H(Y|\mathcal{X}) = H(Y|\mathcal{X}_{\setminus\mathcal{P}'}) = H(Y|\mathcal{X}_{\setminus\mathcal{P}}) &< H(Y|\mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')}) \\ \wedge\ \psi_{\mathcal{P}'}\ \wedge\ \psi_{\mathcal{P}}\Big).&\end{aligned} \tag{3}$$

When $\Psi(Y|\mathcal{X})$ holds, two subsets $\mathcal{P}$ and $\mathcal{P}'$ each reduce the uncertainty of Y by identical amounts—removing either subset from $\mathcal{X}$ leaves the conditional entropy $H(Y|\mathcal{X})$ unchanged, yet removing both increases it. This presents a challenge for feature selection: including both subsets is redundant, but existing methods either include both (Catav et al., 2021; Janssen et al., 2022) or exclude both (Debeer & Strobl, 2020).

6 Approach

In this section, we will present the design of TERC, first discussing a naïve solution and then examining key aspects of the problem, namely perfect conditional redundancy and synergy.

6.1 Overview

To guide the reader, we briefly summarize the structure of the presentation of our approach. The main steps are outlined below:

- **TERC criterion (Section 6.2).** We introduce the mathematical criterion TERC uses to determine whether actions depend on particular observable state variables.
- **Naïve solution (Section 6.3).** We derive a straightforward method for identifying an appropriate representation $\mathcal{X}_*$. This method relies on the assumption that CPMCR does not hold. In practice, this assumption rarely holds, so we refer to it as the naïve solution.
- **Practical method (Section 6.4).** We extend the naïve approach to handle the presence of CPMCR under a weak assumption (Condition 1), yielding a practical and optimal method.

We note that the target representation $\mathcal{X}_*$ is defined as a *member* of a set of subsets, since multiple representations may satisfy the conditions introduced in Section 4.

Practical thresholding (preview). Throughout this section we discuss the criterion in its idealized form, namely whether the quantity $\Phi_{X_i;\mathcal{X}\to A}$ equals zero. In finite samples, exact equality to zero is not a practical test. Our implementation therefore never compares a point estimate to zero directly: instead, as detailed in Section 6.5, we compare each variable's estimated transfer-entropy contribution against a null distribution generated from a known-uninformative random variable with a 95% confidence threshold.

6.2 Design of the Transfer Entropy Redundancy Criterion (TERC)

We now introduce the mathematical details of the Transfer Entropy Redundancy Criterion (TERC). TERC is based on the concept of transfer entropy (TE), as we are interested in how variables *influence* actions using Granger's interpretation (Granger, 1969). Specifically, TERC quantifies the reduction in uncertainty associated with realizations of A when considering the set $\mathcal{X}$ with and without a given variable. If this value is bigger than zero, the actions are said to depend on this state variable, and TERC is verified. More formally, we define TERC as follows:

$$\Phi_{X_i;\mathcal{X}\to A} = H(A|\mathcal{X}_{\setminus X_i}) - H(A|\mathcal{X}) > 0. \tag{4}$$

$\Phi_{X_i;\mathcal{X}\to A}$ describes how actions depend conditionally on observable state variables; therefore, we graphically represent these directed dependencies as Bayesian networks. For a full description of how to estimate the measure defined in Equation 4, refer to Algorithm 2 in Appendix C.

$H(A|\mathcal{X}_{\setminus X_i}) - H(A|\mathcal{X})$ quantifies the reduction in entropy of the realizations of A, when variable X_i is removed from the set $\mathcal{X}$. Positive values of this quantity indicate the actions are dependent on X_i and, therefore, they should be maintained in the state representation. In this case, we describe X_i as satisfying TERC. Conversely, if there is no reduction in the entropy of the values of A, then $\Phi_{X_i;\mathcal{X}\to A} = 0$[5] and the variable does not satisfy TERC. In this case, by definition, we can say that the information provided by variable X_i is either irrelevant or redundant and removing X_i from the state should lead to a policy that can perform inference more efficiently. It follows that we should only add variables to $\mathcal{X}_*$, the subset of maximal information and minimal cardinality, if they satisfy TERC ($\Phi_{X_i;\mathcal{X}\to A} > 0$).

6.3 A Naïve Solution

We now describe a simple application of TERC, our previously defined criterion. To begin, we instantiate an empty set $\mathcal{X}_\Phi$. We then populate this set via the simultaneous addition of variables if their removal from set $\mathcal{X}$ increases the conditional entropy of A. We write this more formally as:

$$\mathcal{X}_\Phi = \{X_i \in \mathcal{X} : \Phi_{X_i;\mathcal{X}\to A} > 0\}. \tag{5}$$

However, in case of CPMCR, the subset $\mathcal{X}_\Phi$ does not satisfy $H(A|\mathcal{X}_\Phi) = H(A|\mathcal{X})$. We formally write this by means of the following lemma:

Lemma 1. *If CPMCR exists in $\mathcal{X}$ (i.e., $\Psi(A|\mathcal{X})$ holds), then $H(A|\mathcal{X}) < H(A|\mathcal{X}_\Phi)$.*

Proof. Refer to Appendix E.

In Lemma 1, we show that, if $\Psi(A|\mathcal{X})$ is verified, $H(A|\mathcal{X}) < H(A|\mathcal{X}_\Phi)$ and, therefore, $\mathcal{X}_\Phi \neq \mathcal{X}_*$. On the contrary, by assuming that there are no cases of CPMCR in $\mathcal{X}$, we can prove the following:

Theorem 1. *If no CPMCR exists in $\mathcal{X}$ (i.e., $\neg\Psi(A|\mathcal{X})$), then the set $\mathcal{X}_\Phi$ defined in Equation 5 satisfies $\mathcal{X}_* = \mathcal{X}_\Phi$.*

Proof. See Appendix F.

However, $\neg\Psi(A|\mathcal{X})$ is not always satisfied.

6.4 TERC in Practice

In the preceding section, we presented a naive application of our criterion TERC, showing that CPMCR cases can prevent us from achieving the goal outlined in the problem statement. In this section, we address this issue by presenting an algorithm—valid under a weak assumption—that applies TERC in a computationally efficient manner. To begin, we motivate the need for our condition before formally introducing it and our algorithm.

[5] This is due to the non-negativity of the measure defined in Equation 4, a property that we prove in Appendix D.

Algorithm 1 A Simple State Variable Selection Method Based on TERC

Input: Observable state and action variables generated by sampling from learning trajectories: $\mathcal{X}$ and A respectively (see Section 3).

Output: The smallest subset of $\mathcal{X}$ that still fully describes the agents actions: $\mathcal{X}_{A_1}$

```
Initialize X_{A_1} = {}
for i = 1 to N do
    if Φ_{X_i;X→A} = 0 then
        X = X \ {X_i}
    else
        X_{A_1} = X_{A_1} ∪ {X_i}
    end
end for
return X_{A_1}
```

Instead of adding features to our selected set simultaneously, suppose we iterate through them randomly and add them one at a time while re-verifying TERC. As a result, if two variables are perfectly redundant, the first one encountered will be removed. Upon re-verifying TERC for the second redundant variable, the removal of the first renders the second non-redundant. Consequently, we negate the issue described in the preceding section. However, if rather than consider two redundant variables we consider two redundant subsets, we have no mechanism that ensures we encounter and remove the smaller of the two subsets. Therefore, we can only be sure that we include the minimum number of variables if both redundant subsets are of equal size. This observation leads us to define the following condition:

Condition 1. *Let $\psi_{\mathcal{P}}$ be as defined in Equation 2 and let $\mathcal{P} \in \mathscr{P}(\mathcal{X})$. We define the condition as follows*

$$\begin{aligned} C_1 = (\forall \mathcal{P} \in \mathscr{P}(\mathcal{X}), \nexists \mathcal{P}' \in \mathscr{P}(\mathcal{X}) : & |\mathcal{P}| \neq |\mathcal{P}'| \quad \& \\ & H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus \mathcal{P}'}) = H(A|\mathcal{X}_{\setminus \mathcal{P}}) < H(A|\mathcal{X}_{\setminus (\mathcal{P},\mathcal{P}')}) \quad \& \\ & \psi_{\mathcal{P}} \quad \& \\ & \psi_{\mathcal{P}'}). \end{aligned} \tag{6}$$

In Appendix B, we show that this condition is satisfied for all of the datasets investigated in this paper. We now more formally describe the approach, which guarantees the derivation of $\mathcal{X}_*$, provided condition C_1 is true. Because the subset derived using this approach requires the use of Algorithm 1 (A_1), it will be labeled $\mathcal{X}_{A_1}$. We write the three steps involved in this approach as follows:

1. Generate trajectories by training an agent using all observable variables as the state $s^t = [x_1^t, x_2^t \dots x_N^t]$, where $s^t \in \mathcal{S}$ and x_i^t is a realization of $X_i \in \mathcal{X}$.
2. Iterate through variables $X_i \in \mathcal{X}$ and remove variables from $\mathcal{X}$ that satisfy $\Phi_{X_i;\mathcal{X}\to A} = 0$, otherwise add them to $\mathcal{X}_{A_1}$.
3. Use this newly designed state for agent training, such that $s_{A_1}^t = [x_1^t, x_2^t \dots x_N^t]$, where $s_{A_1}^t \in \mathcal{S}$ and x_i^t is a realization of $X_i \in \mathcal{X}_{A_1}$.

Notably, the complexity of Algorithm 1 scales linearly in time with respect to the number of features. This makes it preferable to pre-existing techniques that include or exclude features that maximize some function based on the feature added at each step (Brown et al., 2012; Wookey & Konidaris, 2015; Gao et al., 2016; Borboudakis & Tsamardinos, 2019; Tsamardinos et al., 2019; Covert et al., 2023; Bonetti et al., 2024). We now introduce theoretical guarantees for this method.

Theorem 2. *Let $\mathcal{X}_{A_1}$ be the subset generated by Algorithm 1. If Condition 1 (Equation 6) holds, then $\mathcal{X}_* = \mathcal{X}_{A_1}$.*

Proof. See Appendix G.

In practice, this condition is satisfied in many scenarios, including all of the cases studied in the upcoming experimental evaluation.

6.5 Methods for Approximating State Variable Inclusion Conditions

Since we neurally estimate $\Phi_{X_i;\mathcal{X}\to A}$ instead of computing it directly, the estimates are susceptible to noise, which may lead to the inclusion of uninformative variables in our target set. Similarly to Wollstadt et al. (2023), to approximate the condition $\Phi_{X_i;\mathcal{X}\to A} = 0$ for fluctuating estimates, we adopt a null model to which the values of $\Phi_{X_i;\mathcal{X}\to A}$ can be compared. This is done by introducing a random variable, NM, once training is complete. In theory, the actions should have no dependence on NM; therefore, $\Phi_{NM;\mathcal{X}\to A}$ should approach zero under all circumstances. Hence, variables that transfer entropy to actions are expected to deviate from the null model results. We assume normality for both the distribution of the null model and the variables in $\mathcal{X}$. We then use the methods outlined in (Neyman & Pearson, 1933) to determine which $X_i \in \mathcal{X}$ have values of $\Phi_{X_i;\mathcal{X}\to A}$ with a 95% chance of falling outside the range described by the null model. These variables are considered to show a statistically significant deviation from the null model, and therefore satisfy $\Phi_{X_i;\mathcal{X}\to A} > 0$.

In the upcoming experimental evaluation, we will present the upper bound of the null model's 95% confidence interval as a red dashed line. observable state variables whose lower bound exceeds this dashed line are considered to satisfy TERC, and are included in our representation. Otherwise, they are excluded.

7 Experimental Evaluation

We evaluate TERC on synthetic data and RL games of increasing complexity, concluding with the identification of optimal state-history lengths via the Iterated Prisoner's Dilemma.

7.1 General Experimental Settings

Unless otherwise stated, $\Phi_{X_i;\mathcal{X}\to A}$ is averaged over 10 runs while RL environment results are averaged over 5 runs; figures display $\pm 95\%$ confidence intervals with a red dotted line indicating the null-model threshold (Section 6.5). Hyperparameters are tuned via grid search (Appendix I). Across all environments, TERC selects consistent variables across seeds.

7.2 Baselines

We compare against two feature-selection baselines. The first is Ultra Marginal Feature Importance (UMFI) with optimal transport (Janssen et al., 2023), the only existing approach that resolves redundancies while scaling linearly in the number of features. The second is the widely adopted Permutation Importance (PI) algorithm (Breiman, 2001), which is also computationally efficient and is typically insensitive to hyperparameter tuning (Probst et al., 2019). Neither UMFI nor PI, nor the Sequential Knockoffs (SEEK) procedure (Ma et al., 2023), are designed to detect synergistic interactions: their per-feature scoring contrasts cannot, even in principle, recover subsets that are informative only jointly (e.g. XOR-type relationships), which is why we focus on UMFI and PI as the representative practical baselines.

7.3 Experiments on Synthetic Data

7.3.1 Motivation

We first evaluate TERC against the baselines on synthetic data, to obtain controlled and easily interpretable results. Throughout this section (and in subsequent experiments), bar plots compare the per-variable feature-importance scores of TERC, UMFI, and PI. Variables are retained if their scores exceed the null-model threshold (red dashed line) with at least 95% probability.

7.3.2 Synthetic Data Generation

We construct two synthetic datasets exhibiting different forms of CPMCR: the *Four Redundant Variables dataset* (CPMCR between four individual variables) and the *Two Redundant Triplets dataset* (CPMCR between two triplets of variables).

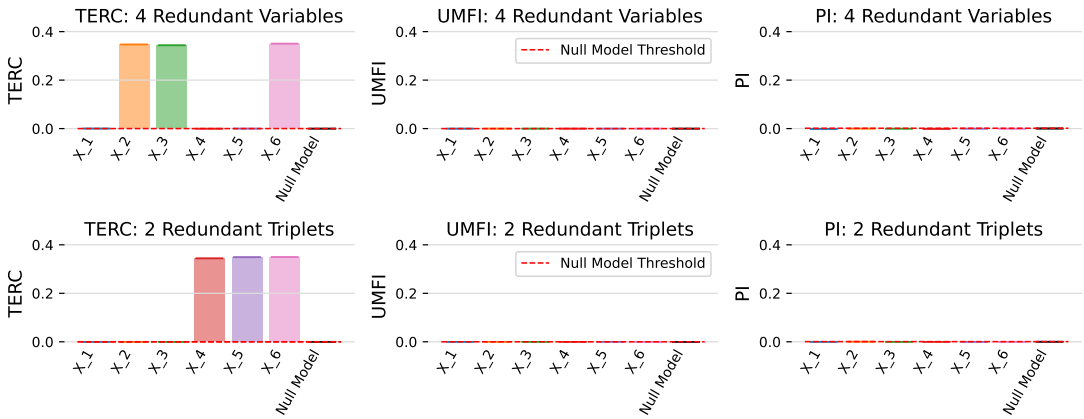

Figure 1: In this graph we illustrate TERC's effectiveness in dealing with complex redundancies and synergies. We plot the values for TERC, PI and UMFI for the *Four Redundant Variables* and the *Two Redundant Triplets* datasets.

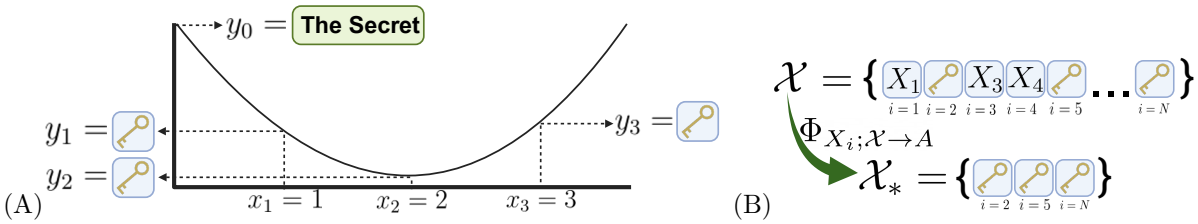

Figure 2: Subfigure (A) illustrates how at least three secret keys are needed to decode a polynomial of order two in Shamir's secure multi-party communication. Subfigure (B) depicts how we use our method of state variable selection to distinguish these three secret-forming keys, from non-secret-forming keys.

Both datasets use six binary variables $X_1, \ldots, X_6$ of length 10,000 and define the target A as the three-way XOR of X_1, X_2, X_3. In the *Four Redundant Variables dataset*, $X_4 \equiv X_5 \equiv X_6 \equiv X_1$, so determining A requires X_2, X_3, and any one of $\{X_1, X_4, X_5, X_6\}$. In the *Two Redundant Triplets dataset*, $X_1 \equiv X_4$, $X_2 \equiv X_5$, $X_3 \equiv X_6$, so CPMCR holds between any two triplets formed by selecting one member from each equivalence class; a feature-selection method should return exactly one such triplet. Both datasets combine redundancy with synergy, which pairwise methods cannot resolve (Section 7.3.1).

7.3.3 Experimental Results

As depicted in Figure 1, for both synthetic datasets, we see that our method is able to successfully identify the minimal set needed to describe the target variable A. This is unlike both UMFI and PI, which fail to identify any variables, because of these methods' inability to resolve the relationships between perfectly redundant variables beyond the pairwise case.

7.4 The Secret Key Game

7.4.1 Motivation

To clearly demonstrate the application and significance of the newly devised technique within the context of RL, we design a new environment, namely the Secret Key Game. The game is inspired by Shamir's secret-sharing protocol developed for secure multiparty communication. This protocol involves dividing a numerical secret into parts known as secret keys and distributing them among group members (Shamir, 1979; Blakley, 1979). In this section, we introduce this protocol before adapting it to an RL game. We model the problem as follows: the agent learns to calculate the secret from a state of N possible secret keys. However,

only three of these keys are actually used to form the secret. The agent's ability to complete its task is consequently impeded until it learns which keys are informative. Therefore, we can leverage our proposed method to exclude non-secret forming keys from the state, resulting in improved game performance. This improvement is attributed to the reduced dimensionality of the state space, in line with Bellman's principles (Bellman & Kalaba, 1959).

We now discuss the secret-sharing protocol that forms the basis for the game. Imagine there is a numerical secret, denoted as y_0, which must be kept secure until all members of a pre-determined group agree to disclose it. To accomplish this, we construct a polynomial function $f(x) = y_0 - ax + bx^2$. We then randomly generate the coefficients a and b in the range $[0, 1]$. By doing so, we obtain a curve with a y-intercept equal to our secret. We then identify three points belonging to this polynomial curve, specifically $(x_1, y_1), (x_2, y_2)$, and (x_3, y_3), and define the y-values of these points (y_1, y_2, and y_3) as our secret keys to be shared among the group members. Figure 2.A visualizes this process, showing a second-order polynomial curve with three secret keys defined by the y-values of three distinct points, while the secret is represented by the y-intercept. The original secret can be derived using polynomial interpolation, but if, and only if, all three secret keys are made available. We adapt this method of multi-party communication into a simple RL game by considering the following question: given you had N keys, of which only three were being used to form y_0, how many incorrect guesses would it take for an RL agent to learn how to correctly calculate the secret?

The Secret Key Game is played as follows: for each iteration of the game, there is a unique new secret, which is divided into three secret keys. These keys are then hidden among $N - K$ decoy keys to form set $\mathcal{X}$, as illustrated in Figure 2.B. In our experiments, $K = 3$. The set $\mathcal{X}$ then serves as the state of the RL agent, considering the following reward function: $r = -|a^t - y_0|$, i.e., the negative absolute difference between the agent's action at time t and the secret. We design the game as described to showcase the effectiveness of the state variable selection method in detecting non-secret forming keys (i.e., keys that are not inputs of the functions used to generate the secret) in $\mathcal{X}$, resulting in their removal to form $\mathcal{X}_*$ as illustrated in Figure 2.B. This enhances the accuracy of the RL agent in approximating the y-intercept of the polynomial function, resulting in improved game performance.

Concretely, with observable state length equal to 25 (three secret keys at randomly chosen indices, e.g., 2, 6, 25, and 22 decoy keys), each iteration draws a fresh observable state s^t consisting of 25 integers in $[0, 10]$. The three secret-key values y_1, y_2, y_3 together with $x \in \{1, 2, 3\}$ define the interpolating polynomial whose y-intercept is the secret y_0. The agent's reward is $r^t = -|a^t - y_0|$ with action space $\mathcal{A} = [-40, 40]$, and the transition function assigns equal probability across states ($T : \mathcal{S} \mapsto \Delta(\mathcal{S})$). The secret-key indices remain fixed across iterations of a given game.

7.4.2 Experimental Settings

In our experiments, we consider states formed of 25 or 50 keys when generating trajectories. By assuming that the y-values of the secret keys are integers in the range [0..10]; the resulting secret is also an integer within the range of $y_0 \in [-40..40]$. Therefore, to ensure that the agent action space is the same as the range of the secret, we employ an RL algorithm with an 80-dimensional discrete action space $\mathcal{A} = [-40..40]$. Due to the large observable state spaces, we use a one-step temporal difference learning Actor-Critic algorithm (Sutton & Barto, 2018), as described in 6.4. Finally, as far as the selection of the hyperparameters for the RL algorithm is concerned, we refer to Appendix I.

7.4.3 Experimental Results

Figure 3.A plots the values of $\Phi_{X_i;\mathcal{X}\to A}$ that are calculated when aiming to verify TERC. Furthermore, in this plot, we also display values of UMFI and PI for each key. In Figure 3.A, there are 25 potential keys in the state, of which only three form the secret, which we label as secret keys, while the remainder have been labeled as decoy keys. According to the results of TERC in Figure 3.A, the actions were only conditionally dependent on the three secret forming keys, which were randomly selected at the start of the game as keys $2, 6$ and 25. We represent the results using a Bayesian network as illustrated in Figure 3.B. UMFI similarly leads to a set consisting of all three secret keys when applied to this simple game. However, PI can only identify two out of three secret keys. We observe similar results in 3.A when the observable state is composed of 50 keys.

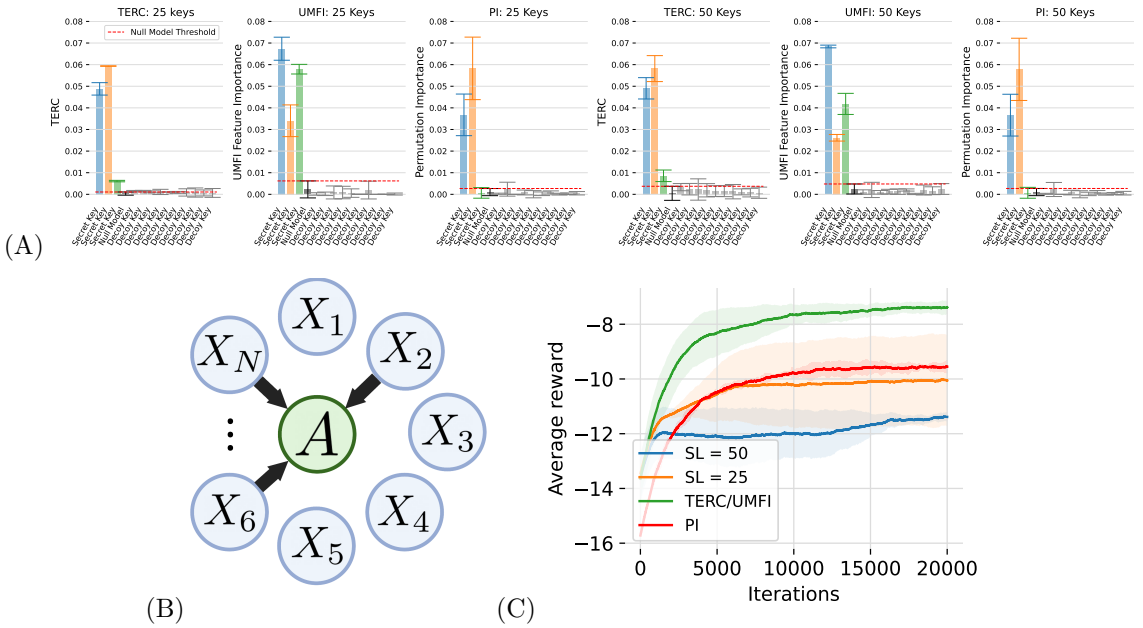

Figure 3: Subfigure (A) depicts the final feature importance values for each key in the secret game when using TERC, UMFI or PI. Subfigure (B) depicts the Bayesian network representation of TERC in the Secret Key Game if the secret keys were at index two, six and N. Finally, graph (C) represents how the agent training efficiency varied as a function of state length.

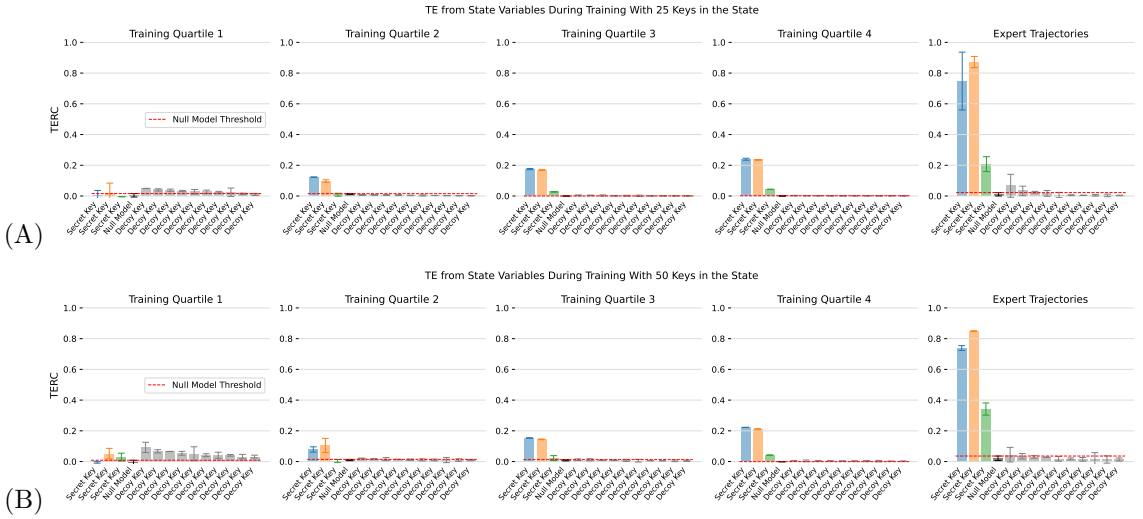

Figure 4: Graphs (A) and (B) show the final values computed when verifying TERC as evaluated during different training quartiles, for 25 and 50 potential secret keys. Similarly to Figure 3, we have included only TERC values for the 10 decoy keys that transferred the most entropy to the actions.

Figure 3.C shows how failing to implement state variable selection can deteriorate the agent performance while playing the Secret Key Game. We observe that the state designed using TERC/UMFI achieves the greatest cumulative reward. Beyond training performance, the reduced state representation also improves deployment efficiency. Table 1 reports inference times for policies trained on full versus TERC-selected states. The substantial dimensionality reduction (88–94%) yields speedups of up to 2.6× on CPU and 1.9× on GPU for SKG-50.

Table 1: Inference speedup for the Secret Key Game on CPU caused by using the state selected by TERC, averaged over 5 seeds (10,000 inference calls each). B denotes batch size.

Environment	B=1	B=64
SKG-25 (25 → 3)	1.43×	1.47×
SKG-50 (50 → 3)	2.60×	2.20×

Figure 4.A and 4.B depict how $\Phi_{X_i;\mathcal{X}\to A}$ evolves during different stages of training when the state comprises either 25 or 50 keys. As the agent learns, it receives higher rewards by enhancing its ability to decipher the secret. This leads to the agent's actions increasingly depending on the secret keys, as shown in Figures 4.A and 4.B.

7.5 Gym Physics Environments

7.5.1 Motivation

In order to show the potential of the proposed method for physics-based RL problems, we use three environments provided by OpenAI Gym (Brockman et al., 2016), namely Cart Pole, Lunar Lander, and Pendulum.

Following (Grooten et al., 2023), we 'dope' the environment-provided state with a set of random variables V_{rand_i}. For example, in Cart Pole[6] the observable state $s^t = [x^t, \dot{x}^t, \theta^t_{pole}, \dot{\theta}^t_{pole}]$ is augmented with three noise variables to give $\mathcal{X} = \{X, \dot{X}, \Theta_{pole}, \dot{\Theta}_{pole}, V_{rand_1}, V_{rand_2}, V_{rand_3}\}$. We train RL agents on the doped observable state until convergence, then verify that TERC removes V_{rand_i} to recover $\mathcal{X}_*$, and that retraining on $\mathcal{X}_*$ improves learning efficiency. Examining $\Phi_{X_i;\mathcal{X}\to A}$ across training stages additionally reveals distinct phases in the behavioral dynamics of the agent, highlighting TERC's potential for interpretability.

7.5.2 Experimental Settings

For the environments with discrete action and state spaces, namely, Cart Pole and Lunar Lander, we train the agents using a one-step temporal difference Actor-Critic architecture (Sutton & Barto, 2018). Instead, for the continuous action space environment we take into consideration, Pendulum, we use PPO (Schulman et al., 2017). For more details on these methods, including hyperparameters, please refer to Appendix I.

7.5.3 Experimental results

The $\Phi_{X_i;\mathcal{X}\to A}$ and UMFI values of the completely random variables (V_{rand_i}), as depicted in Figure 5, fall within the range of the null model. As a result, these random variables are excluded from the set $\mathcal{X}$, to form $\mathcal{X}_*$. Consequently, this updated set contains only the original variables. PI instead fails to detect all the informative variables. Training the agent using the sets identified using TERC, UMFI and PI leads to the learning curves seen on the right-hand side of Figure 5. These graphs show how training can be sped up when using the optimal set of variables as identified using TERC and UMFI. Table 2 reports that these reduced representations also yield modest inference speedups (1.05–1.33× on CPU), with gains scaling proportionally to the degree of dimensionality reduction.

Table 2: Inference speedup for Gym environments on CPU, averaged over 5 seeds (10,000 inference calls each). B denotes batch size.

Environment	B=1	B=64
CartPole (7 → 4)	1.16×	1.33×
LunarLander (11 → 8)	1.10×	1.25×
Pendulum (6 → 3)	1.05×	1.03×

[6] `https://gymnasium.farama.org/environments/classic_control/cart_pole/`.

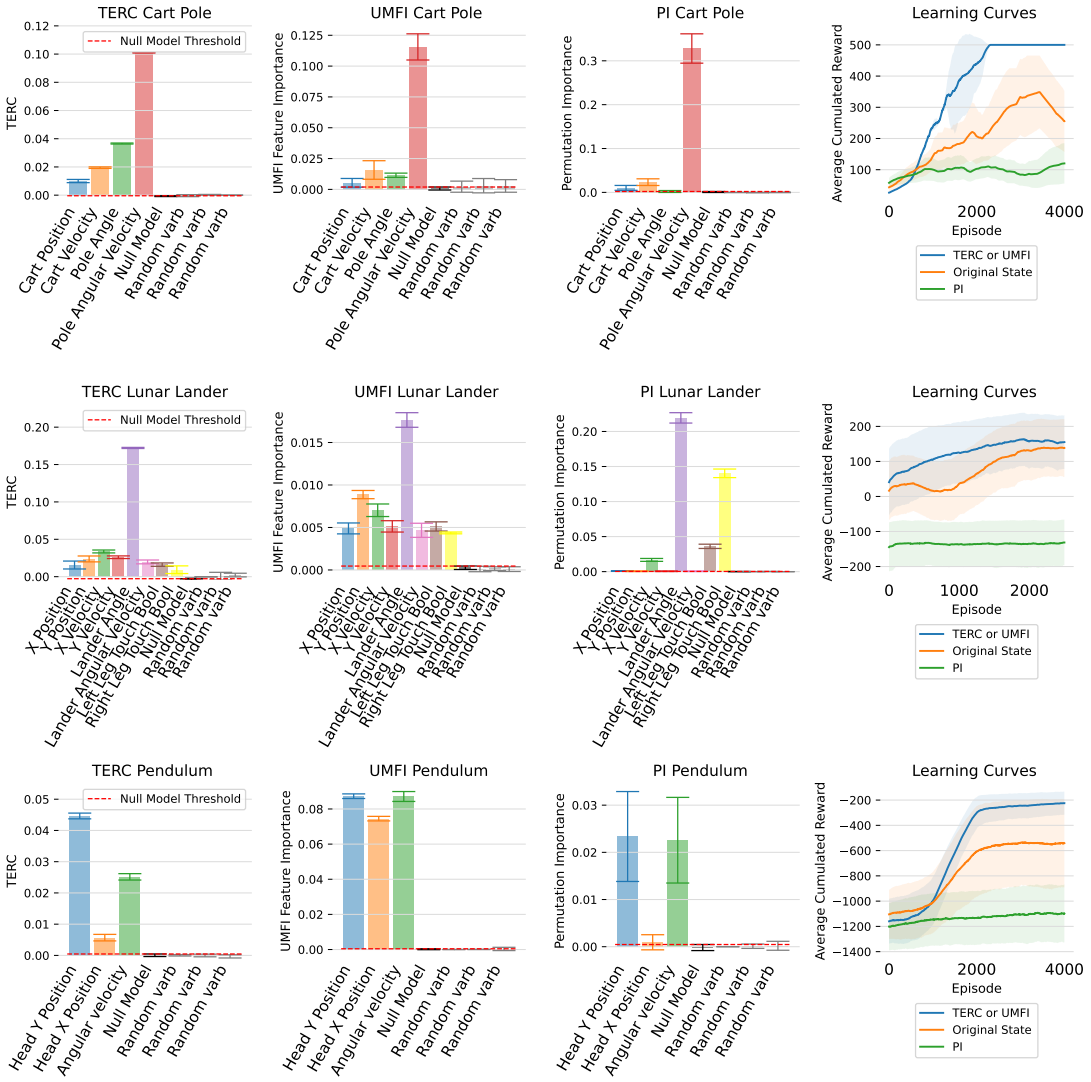

Figure 5: This subfigure depicts the final values obtained for $\Phi_{X_i;\mathcal{X}\to A}$, UMFI, and PI for Cart Pole, Lunar Lander, and Pendulum. On the right-hand side, we illustrate how failing to remove these random variables from the state of the Cart Pole playing agent degrades the performance of the game.

Now, we examine how $\Phi_{X_i;\mathcal{X}\to A}$ changes as the agent trains, demonstrating TERC's potential use for model interpretability. Figure 6 illustrates these changes as the agent learns the cart-pole game. Initially, the agent learns to maintain the pole's balance by moving in one direction according to the sign of the angular velocity until it reaches the boundary. As a result, during the first quarter of the training, we observe that the agent's actions depend only on the angular velocity. It is possible to observe that the agent tries to avoid the environment's edge during the second quartile of training, which leads to an increase of the value of $\Phi_{X_i;\mathcal{X}\to A}$ for the variables associated with cart velocity and x-position. As the agent moves into the fourth quartile, it is increasingly able to remain within the central area of the environment, reducing the need for correcting the cart position; as a result, its actions exhibit less dependence on the cart position variable. In the final graph of Figure 6 (right), the agent has perfected its ability to stay in the center of the environment, making only minor movements to re-adjust the angle of the pole if it deviates from the upright position. This shift in strategy leads to actions that show less dependence on the angular velocity, velocity, and position while exhibiting a stronger dependence on the angle.

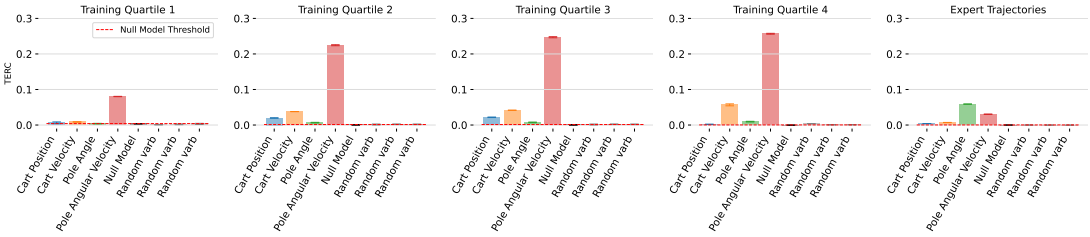

Figure 6: Entropy transferred from the observable state variables to the actions during different phases of training in the Cart Pole environment.

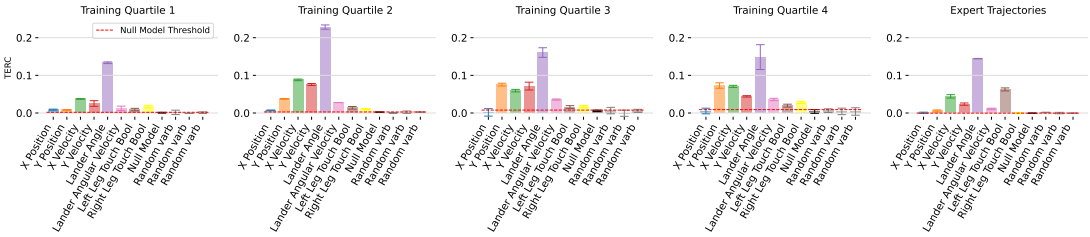

Figure 7: Entropy transferred from the observable state variables to the actions during different phases of training in the Lunar Lander environment.

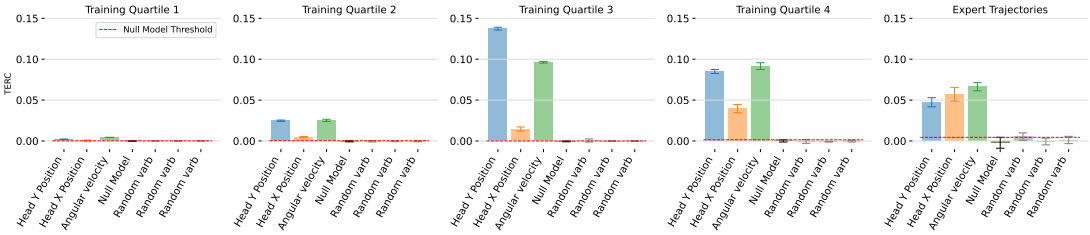

Figure 8: Entropy transferred from the observable state variables to the actions during different phases of training in the Pendulum environment.

Figure 7 instead shows how an RL agent's strategy evolves while learning to play Lunar Lander. Initially, during the first quarter of training, the agent discovers the necessity of taking actions that are dependent on the lander's angle to maintain its upright position. After achieving this, the agent begins to learn to fly around the environment, leading to its actions becoming more dependent upon the speed and position variables throughout the second and third quarters of training. Finally, throughout the trajectories collected from the last quarter of training and the expert trajectories, the agent masters landing within the target area, which results in the actions depending more on the Boolean variable that describes whether the lander's legs have touched down successfully.

Figure 8 demonstrates the changes in $\Phi_{X_i;\mathcal{X}\to A}$ estimates as the agent learns how to play the Pendulum game. During the first quartile, the agent's actions are primarily stochastic due to the lack of a learned strategy, resulting in negligible $\Phi_{X_i;\mathcal{X}\to A}$ values across all variables. In the second and third quartiles, the agent learns to swing the pendulum to the upright position, making its actions increasingly dependent on the y-position of the pendulum head and the angular velocity, as depicted in Figures 8.B and 8.C. Once the agent has learned to swing the pendulum upright, it then learns to keep it there by making minor side-to-side adjustments. This is reflected in the values of $\Phi_{X_i;\mathcal{X}\to A}$ reported in the third and fourth quartiles, where we observe a rising dependency on the x position value of the pendulum head.

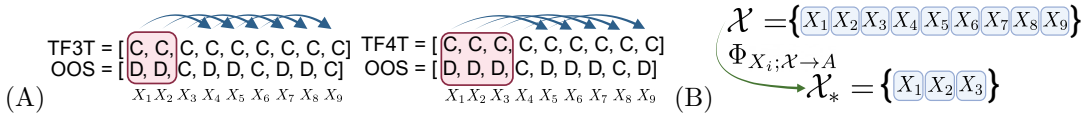

Figure 9: In Subfigure (A) we depict the moves made by both players when playing optimally against TF3T and TF4T strategies in the Iterated Prisoner's Dilemma. The colored boxes indicate the minimum history needed to learn this optimal strategy, while the colored arrows indicate one period of a repeating action sequence, where C stands for cooperative moves, while D represents defective ones. In Subfigure (B), we illustrate schematically how our measure is used to determine the optimal state history length of three, out of a maximum of nine, when facing a TF4T opponent.

7.6 Tit-For-N-Tats Strategy in the Iterated Prisoner's Dilemma

7.6.1 Motivation

Our final environment is the Iterated Prisoner's Dilemma (Axelrod & Hamilton, 1981), which we use to demonstrate that TERC can discover optimal history lengths when temporally extended states are required. The state $s^t \in \mathcal{S}$ consists of the past l actions of both the player and their opponent, with each action $a_i^t \in \{C, D\}$ denoting cooperation or defection (Anastassacos et al., 2020). The reward matrix is given in Appendix I.6. When both contestants are maximizing reward under this matrix, $l = 1$ suffices for optimality (Press & Dyson, 2012); however, against other opponent policies a longer history may be required. A canonical case is Tit-For-N-Tats (TFNT): the opponent defects only after N consecutive defections and otherwise cooperates. The optimal counter-strategy takes $N-1$ defective actions before cooperating, which requires a minimum history length of $N-1$ (red boxes in Figure 9.A, showing the optimal opponent strategies against TF3T and TF4T). We demonstrate below that TERC reliably identifies this optimal history length.

7.6.2 Experimental Settings

Let l denote the maximum history length of interest. We train a tabular Q-learning agent (hyperparameters in Appendix I) on a state of history length l against a TFNT opponent, and form $\mathcal{X} = \{X_1 \dots X_l\}$ from the resulting action-pair variables (Figure 9.A). We set $L = 9$ to cover strategies up to Tit-For-10-Tats. Because the optimal counter-strategy is cyclic with period N, the realizations of X_i and X_{i+N} coincide, yielding CPMCR both between the pair $\{X_i, X_{i+N}\}$ ($\Psi_{X_i,X_{i+N}}(A|\mathcal{X})$) and between the full-period subsets $\mathcal{P} = \{X_i, X_{i+1}, \dots, X_{i+N-1}\}$ and $\mathcal{P}' = \{X_{i+N}, X_{i+N+1}, \dots, X_{i+2N-1}\}$ (colored arrows in Figure 9.A). This is the only non-synthetic experiment in which CPMCR arises in our empirical evaluation. Network hyperparameters for Algorithm 1 are given in Appendix I.

7.6.3 Experimental Results

In Figure 10.A, we show the values of $\Phi_{X_i;\mathcal{X}\to A}$, UMFI and PI when applied to the trajectories generated in the case of a Q-learner play against a TFNT opponent. In these datasets we do not have simple pairwise redundancy and, for this reason, the assumptions in (Janssen et al., 2023) are not verified. Therefore, this method is not able to identify the correct set of variables; PI is similarly ineffective. Only TERC has the ability to identify the correct set of variables. We represent the values in the TF4T case pictorially as the Bayesian network shown in Figure 10.A. This is further reflected in the agent performance results reported in Figure 10.C, where we plot the normalized cumulative reward achieved over the last 1000 iterations of the game, as a function of the length of the state history. In Figure 10.C we report the percentage of times the agent plays optimally during the last 1000 iterations of training. Agents trained with a history length smaller than $N-1$ converge to suboptimal strategies. Consequently, we observe that only state lengths equal to 9 achieve optimal cumulative rewards against the TF10T opponent, whereas state lengths of 2 play optimally against the TF3T opponent. The variables that satisfy TERC in Figure 10.C can therefore be interpreted as indicating the minimum subset of $\mathcal{X}$ required to play optimally, corroborating the theoretical observations reported earlier. The detailed learning dynamics are reported in the Appendix H. TERC selects

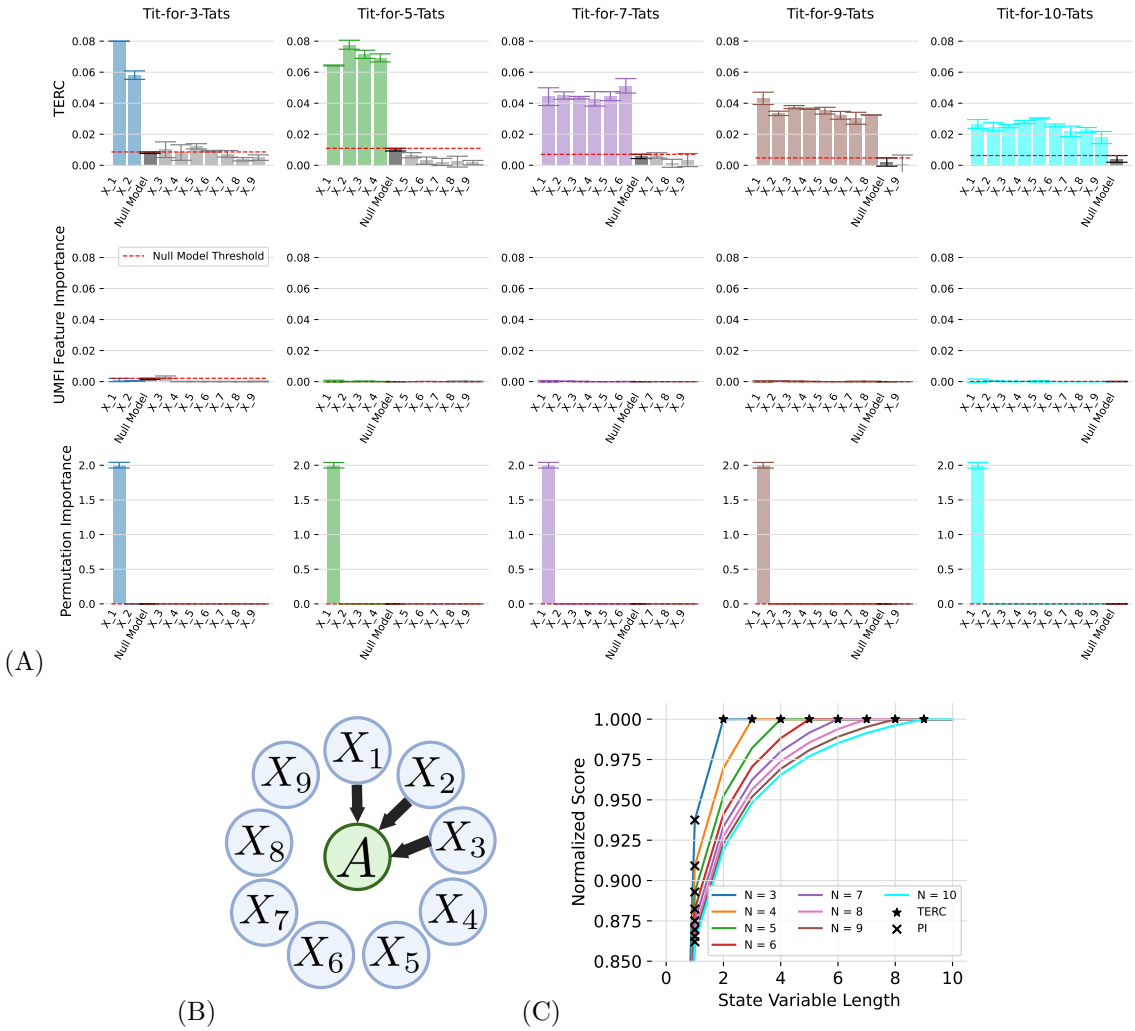

Figure 10: In Subfigure (A) we compare $\Phi_{X_i;\mathcal{X}\to A}$, UMFI and PI, values for all state variables from $l = 1$ to $l = 9$, when playing optimally against a TFNT opponent. For clarity of presentation, we only present the results with the values of N displayed above. These results were representative of how our algorithm performed for all TFNT opponents. Subfigure (B) shows the resulting Bayesian network representation of how the actions depend on observable state variables against a TF4T opponent. Subfigure (C) shows how the change of the state size affects an agent's performance against a TFNT strategy. We use stars to denote the optimal state length highlighted by our method, whereas we use crosses to represent the state length generated using PI. Finally, we omit the results for UMFI in Subfigure (C) for clarity.

consistent state variables across seeds. Inference improvements are not reported because we employ tabular Q-learning.

8 Conclusions

In this article, we have introduced a novel information-theoretic methodology for state variable selection in RL, based on TERC, a criterion that can be used to verify the existence of dependencies between observable state variables and the agent's actions. The objective is the definition of the minimal subset of state variables from which an agent could learn optimal policies. We have defined this optimal subset as one whose realizations reduce the entropy of the actions identically to those of the set of all observable state variables.

We have demonstrated that, assuming CPMCR is not verified, a naïve application of TERC is sufficient to identify the minimal set (Theorem 1). However, this assumption is not universally satisfied. Consequently, we have also presented a solution in case of CPMCR. We have presented an implementation, underpinned by theoretical results, which guarantees the derivation of the minimal set with a weak assumption (Theorem 2). In general, TERC represents a novel approach toward feature selection that scales linearly in time without relying on the restricting assumption of having only pairwise dependencies in the set of variables. Furthermore, even if C_1 is not satisfied, we still select an information-theoretically optimal set; but we cannot ensure it is also minimal.

Finally, we have implemented our method using a neural estimator and we have then evaluated TERC using synthetic data and across diverse environments representative of various RL problem classes. In particular, we have demonstrated that our proposed method consistently identifies the optimal set of variables in presence of redundant and synergistic relationships. Moreover, we have shown that the resulting compact state representations yield concrete deployment benefits, with inference speedups of up to 2.6× when dimensionality reduction is substantial. These efficiency gains are measured when a policy is retrained or used for inference on the TERC-selected subset; they do not account for the cost of the initial full-state training run used to generate the trajectories on which TERC is applied.

References

Nicolas Anastassacos, Stephen Hailes, and Mirco Musolesi. Partner selection for the emergence of cooperation in multi-agent systems using reinforcement learning. In *AAAI'20*, 2020.

Daniel W. Apley and Jingyu Zhu. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society Series B*, 82(4):1059–1086, 2020.

Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981.

Philip Bachman, R. Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *NeurIPS'19*, 2019.

Roberto Battiti. Using mutual information for selecting features in supervised neural net learning. *IEEE Transactions on Neural Networks*, 5(4):537–550, 1994.

Ishmael Belghazi, Sai Rajeswar, Aristide Baratin, R. Devon Hjelm, and Aaron Courville. Mine: Mutual information neural estimation. In *ICML'18*, 2018.

Anthony Bell and Terrence Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(12):1129–59, 1995.

Richard Bellman and Robert Kalaba. A mathematical theory of adaptive control processes. *Proceedings of the National Academy of Sciences*, 45(8):1288–1290, 1959.

George Blakley. Safeguarding cryptographic keys. In *MaRK'79*, 1979.

Paolo Bonetti, Alberto Maria Metelli, and Marcello Restelli. Causal feature selection via transfer entropy. In *IJCNN'24*, 2024.

Giorgos Borboudakis and Ioannis Tsamardinos. Forward-backward selection with early dropping. *Journal of Machine Learning Research*, 20(1):276–314, 2019.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. In *arXiv:1606.01540*, 2016.

Gavin Brown, Adam Pocock, Ming-Jie Zhao, and Mikel Luján. Conditional likelihood maximisation: A unifying framework for information theoretic feature selection. *Journal of Machine Learning Research*, 13 (2):27–66, 2012.

Amnon Catav, Boyang Fu, Yazeed Zoabi, Ahuva Libi Weiss Meilik, Noam Shomron, Jason Ernst, Sriram Sankararaman, and Ran Gilad-Bachrach. Marginal contribution feature importance - an axiomatic approach for explaining data. In *ICML'21*, 2021.

Hugh Chen, Joseph D. Janizek, Scott Lundberg, and Su-In Lee. True to the model or true to the data? In *arXiv:2006.16234*, 2020.

Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. Learning to explain: An information-theoretic perspective on model interpretation. In *ICML'18*, 2018.

Shay Cohen, Gideon Dror, and Eytan Ruppin. Feature selection via coalitional game theory. *Neural Computation*, 19(8):1939–61, 2007.

Ian C. Covert, Scott Lundberg, and Su-In Lee. Understanding global feature contributions with additive importance measures. In *NeurIPS'20*, 2020.

Ian Connick Covert, Wei Qiu, Mingyu Lu, Na Yoon Kim, Nathan J. White, and Su-In Lee. Learning to maximize mutual information for dynamic feature selection. In *ICML'23*, 2023.

Thomas Dean and Robert Givan. Model Minimization in Markov Decision Processes. In *AAAI'97*, 1997.

Dries Debeer and Carolin Strobl. Conditional permutation importance revisited. *BMC Bioinformatics*, 21 (1):1–30, 07 2020.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *ACL'19*, 2019.

Siegmund Duell, Alexander Hans, and Steffen Udluft. The markov decision process extraction network. In *ESANN'10*, 2010.

Chelsea Finn, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, and Pieter Abbeel. Deep spatial autoencoders for visuomotor learning. In *ICRA'16*, 2016.

Christopher Frye, Colin Rowat, and Ilya Feige. Asymmetric Shapley Values: Incorporating Causal Knowledge into Model-Agnostic Explainability. In *NeurIPS'20*, 2020.

Shuyang Gao, Greg Ver Steeg, and Aram Galstyan. Variational information maximization for feature selection. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (eds.), *NeurIPS'16*, 2016.

Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc Bellemare. Deepmdp: Learning continuous latent space models for representation learning. In *ICML'19*, 2019.

Robert Givan, Thomas Dean, and Matthew Greig. Equivalence notions and model minimization in Markov decision processes. *Artificial Intelligence*, 147(1):163–223, 2003.

Clive William John Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.

Bram Grooten, Ghada Sokar, Shibhansh Dohare, Elena Mocanu, Matthew E Taylor, Mykola Pechenizkiy, and Decebal Constantin Mocanu. Automatic noise filtering with dynamic sparse training in deep reinforcement learning. In *AAMAS'23*, 2023.

Botao Hao, Yaqi Duan, Tor Lattimore, Csaba Szepesvári, and Mengdi Wang. Sparse feature selection makes batch reinforcement learning more sample efficient. In *ICML'21*, 2021.

Daniel Hein, Steffen Udluft, and Thomas A Runkler. Generating interpretable fuzzy controllers using particle swarm optimization and genetic programming. In *GECCO '18*, 2018.

R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR'19*, 2019.

Robert L. Smith James C. Bean, John R. Birge. Aggregation in dynamic programming. *Operations Research*, 35(2):215–220, 1987.

Alexander Janssen, Mark Hoogendoorn, Marjon H. Cnossen, Ron A. A. Mathôt, and the OPTI-CLOT Study Group and SYMPHONY Consortium. Application of SHAP values for inferring the optimal functional form of covariates in pharmacokinetic modeling. *CPT: Pharmacometrics & Systems Pharmacology*, 11(8): 1100–1110, 2022.

Joseph Janssen, Vincent Guan, and Elina Robeva. Ultra-marginal feature importance: Learning from data with causal guarantees. In *AISTATS'23*, pp. 10782–10814, 2023.

Nicholas K. Jong and Peter Stone. State abstraction discovery from irrelevant state variables. In *IJCAI'05*, 2005.

Alon Keinan, Claus C. Hilgetag, Isaac Meilijson, and Eytan Ruppin. Causal localization of neural function: the Shapley value method. *Neurocomputing*, 58-60(3):215–222, 2004.

Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E*, 69(6):066138, 2004.

Elizabeth Kumar, Suresh Venkatasubramanian, Carlos Scheidegger, and Sorelle A. Friedler. Problems with Shapley-Value-Based Explanations as Feature Importance Measures. In *ICML'20*, 2020.

Yongchan Kwon and James Y. Zou. WeightedSHAP: analyzing and improving Shapley based feature attributions. In *NeurIPS'22*, 2022.

Michael Laskin, Aravind Srinivas, and Pieter Abbeel. CURL: Contrastive unsupervised representations for reinforcement learning. In *ICML'20*, 2020.

Timothée Lesort, Natalia Díaz-Rodríguez, Jean-Franois Goudou, and David Filliat. State representation learning for control: An overview. *Neural Networks*, 108:379–392, 2018.

Lihong Li, Thomas Walsh, and Michael Littman. Towards a Unified Theory of State Abstraction for MDPs. In *ISAIM'06*, 2006.

Ralph Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988.

Feng Liu, Ruiming Tang, Xutao Li, Weinan Zhang, Yunming Ye, Haokun Chen, Huifeng Guo, Yuzhou Zhang, and Xiuqiang He. State representation modeling for deep reinforcement learning based recommendation. *Knowledge-Based Systems*, 205:106170, 2020.

Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *NeurIPS'17*, 2017.

Tao Ma, Hengrui Cai, Zhengling Qi, Chengchun Shi, and Eric B. Laber. Sequential knockoffs for variable selection in reinforcement learning. *arXiv preprint arXiv:2303.14281*, 2023.

Andrew Kachites McCallum and Dana Ballard. *Reinforcement Learning with Selective Perception and Hidden State*. PhD thesis, The University of Rochester, 1996.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(6):529–533, 2015.

Jerzy Neyman and Egon Pearson. The testing of statistical hypotheses in relation to probabilities a priori. *Mathematical Proceedings of the Cambridge Philosophical Society*, 29(4):492–510, 1933.

Jennifer Ortiz, Magdalena Balazinska, Johannes Gehrke, and S. Sathiya Keerthi. Learning state representations for query optimization with deep reinforcement learning. In *DEEM'18*, 2018.

Xinlei Pan, Yurong You, Ziyan Wang, and Cewu Lu. Virtual to real reinforcement learning for autonomous driving. In *arXiv:1704.03952*, 2017.

David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. The carbon footprint of machine learning training will plateau, then shrink. *Computer*, 55(7):18–28, 2022.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Hanchuan Peng, Fuhui Long, and Chris Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238, 2005.

Gregory Plumb, Denali Molitor, and Ameet Talwalkar. Model agnostic supervised local explanations. In *NeurIPS'18*, 2018.

Ben Poole, Sherjil Ozair, Aaron Oord, Alexander Alemi, and George Tucker. On variational bounds of mutual information. In *ICML'19*, 2019.

William H Press and Freeman J. Dyson. Iterated Prisoner's Dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences*, 109(26):10409–10413, 2012.

Philipp Probst, Anne-Laure Boulesteix, and Bernd Bischl. Tunability: Importance of hyperparameters of machine learning algorithms. *Journal of Machine Learning Research*, 20(1):1934–1965, 2019.

Balaraman Ravindran and Andrew G. Barto. SMDP Homomorphisms: An Algebraic Approach to Abstraction in Semi-Markov Decision Processes. In *IJCAI'03*, 2003.

Daniele Reda, Tianxin Tao, and Michiel van de Panne. Learning to locomote: Understanding how environment design matters for deep reinforcement learning. In *MIG'20*, 2020.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *SIGKDD'16*, 2016.

Anton Maximilian Schäfer and Steffen Udluft. Solving partially observable reinforcement learning problems with recurrent neural networks. In *Workshop Proceedings of the European Conference on Machine Learning*, 2005.

Anton Maximilian Schäfer, Daniel Schneegass, Volkmar Sterzing, and Steffen Udluft. A neural reinforcement learning approach to gas turbine control. In *IJCNN'07*, 2007.

Thomas Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461–464, 2000.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. In *arXiv:1707.06347*, 2017.

Max Schwarzer, Ankesh Anand, Rishab Goel, R. Devon Hjelm, Aaron Courville, and Philip Bachman. Data-efficient reinforcement learning with self-predictive representations. In *ICLR'20*, 2020.

Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11): 612–613, 1979.

Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3): 379–423, 1948.

Lloyd Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–318, 1953.

Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *ICML'17*, 2017.

Jiaming Song and Stefano Ermon. Understanding the limitations of variational mutual information estimators. In *ICLR'20*, 2020.

Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. In *ICML'21*, 2021.

Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. In *arXiv:1906.05743*, 2019.

Mukund Sundararajan and Amir Najmi. The Many Shapley Values for Model Explanation. In *ICML'20*, 2020.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, 2018.

Mauricio Tesmer and Pablo A. Estévez. AMIFS: Adaptive feature selection by using mutual information. *IJCNN'04*, 2004.

Ioannis Tsamardinos, Giorgos Borboudakis, Pavlos Katsogridakis, Polyvios Pratikakis, and Vassilis Christophides. A greedy feature selection algorithm for Big Data of high dimensionality. *Machine Learning*, 108:149–202, 2019.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. In *arXiv:1807.03748*, 2019.

Paul L. Williams and Randall D. Beer. Nonnegative decomposition of multivariate information. In *arXiv:1004.2515*, 2010.

Patricia Wollstadt, Sebastian Schmitt, and Michael Wibral. A rigorous information-theoretic definition of redundancy and relevancy in feature selection based on (partial) information decomposition. *Journal of Machine Learning Research*, 24(131):1–44, 2023.

Dean S. Wookey and George D. Konidaris. Regularized feature selection in reinforcement learning. *Machine Learning*, 100(2-3):655–676, 2015.

Lin Yang and Mengdi Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *ICML'20*, 2020.

A Terminology and Notation Summary

Terminology. Throughout this paper, we distinguish between several closely related but conceptually distinct notions, defined as follows.

Observable State: $\mathcal{X} = \{X_1, \ldots, X_N\}$ denotes the set of all state variables observable at a given time-step t.

State: the concrete vector used as input to the policy at time t, $s^t \in \mathcal{S}$, built from (a subset of) the observables, possibly stacked over a history window.

Markovian state: a state representation such that the transition and reward distributions satisfy the Markov property, i.e. conditional on the current state the future is independent of the past.

Quasi-Markovian state: a compressed representation (e.g. produced by an RNN over a window of observables) that approximates a Markovian state for a partially observable problem (Schäfer & Udluft, 2005; Schäfer et al., 2007).

Minimal Markovian state: a Markovian state of smallest cardinality among those that preserve $H(A \mid \mathcal{X})$ for the trained policy; this is the target representation $\mathcal{X}_*$ derived by TERC.

Notation summary. Table 3 summarizes the main symbols used throughout the paper.

Table 3: Notation summary.

Symbol	Meaning
$\mathcal{X} = \{X_1, \ldots, X_N\}$	Set of all observable state variables.
X_i, x_i	Observable random variable and its realization.
A, a^t	Action random variable and realization at time t.
$s^t \in \mathcal{S}$	State vector at time t, $s^t = [x_1^t, \ldots, x_N^t]$.
$\mathcal{P}, \mathcal{P}' \subseteq \mathcal{X}$	Subsets of observables.
$\mathscr{P}(\mathcal{X})$	Power set of $\mathcal{X}$.
$\mathcal{X}_{\setminus \mathcal{P}}$	Set difference $\mathcal{X} \setminus \mathcal{P}$.
$\mathcal{X}_{\setminus (\mathcal{P}, \mathcal{P}')}$	Shorthand for $\mathcal{X} \setminus (\mathcal{P} \cup \mathcal{P}')$.
$H(\cdot)$, $H(\cdot \mid \cdot)$	Shannon (conditional) entropy.
$TE_{Y \to X}$	Transfer entropy from Y to X.
$\Phi_{X_i; \mathcal{X} \to A}$	TERC measure for X_i (Eq. 4).
$\Psi(A \mid \mathcal{X})$	CPMCR condition (Eq. 3).
C_1	Condition 1 (Eq. 6).
$\mathcal{X}_\Phi$	Naïve TERC-selected subset (Eq. 5).
$\mathcal{X}_{A_1}$	Subset returned by Algorithm 1.
$\mathcal{X}_*$	Target minimal subset (Eq. 1).
NM	Null-model random variable (Sec. 6.5).

B Empirical Verification of Condition 1

In this section, we empirically verify that Condition 1 (Equation 6) holds across all environments considered in this paper. Recall that Condition 1 states that no two different-sized subsets of the data variables should provide perfectly redundant information about the actions. If this condition is violated, TERC's Algorithm 1 may fail to identify the minimal state representation.

B.1 Methodology

For each environment, we verify Condition 1 as follows. For all disjoint subset pairs $(\mathcal{P}_1, \mathcal{P}_2)$ where $\mathcal{P}_1, \mathcal{P}_2 \subseteq \mathcal{X}$ and $|\mathcal{P}_1| \neq |\mathcal{P}_2|$, we compute the redundancy:

$$R(\mathcal{P}_1; \mathcal{P}_2 \to A) = I(\mathcal{P}_1; A) + I(\mathcal{P}_2; A) - I(\mathcal{P}_1, \mathcal{P}_2; A), \tag{7}$$

where $I(\cdot; A)$ denotes mutual information with the action variable (Williams & Beer, 2010). Condition 1 is violated if $R(\mathcal{P}_1; \mathcal{P}_2 \to A) = H(A)$ for any subset pair, where $H(A)$ is the entropy of the action distribution.

We collect 10,000 state-action pairs from trained policies for each environment and estimate mutual information using the KSG estimator (Kraskov et al., 2004). For environments with discrete state and action spaces (Secret Key Game, IPD), we compute mutual information directly from empirical distributions.

B.2 Results

Figure 11 presents the redundancy values for the 10 most redundant subset pairs in each environment. The dashed line indicates $H(A)$, the threshold above which Condition 1 would be violated (shaded region). Across all environments,including the Secret Key Game (Figure 11.A), Gym physics environments (Figure 11.B), and the Iterated Prisoner's Dilemma against TFNT opponents (Figure 11.C), no subset pair exceeds this threshold. This confirms that Condition 1 is satisfied in all experimental settings, validating the theoretical guarantees of Algorithm 1.

C Algorithm for the Estimation of Transfer Entropy Based Measure

In this section, we outline an algorithm for the estimation of the transfer entropy based measure defined in Equation 4.

Algorithm 3 Estimation of $\Phi_{X_i;\mathcal{X}\to A}$.

Input: Training trajectories $\tau = (s^1, a^1, s^2, a^2 \dots s^T, a^T) = \{\mathcal{X}, A\}$, where $a^1, a^2 \in A$ and $s^1, s^2 \in \mathcal{X}$. Target variable X_i.
Output: $H(A|\mathcal{X}_{\setminus X_i}) - H(A|\mathcal{X})$

1: Initialize weights for θ and $\theta_{\setminus X_i}$
2: **for** 1 to N **do**
3: Draw mini batch samples of length b from the joint distribution of the actions and the state with all possible variables included $p_{A,\mathcal{X}} \sim (a^{t_1}, x_1^{t_1}, x_i^{t_1} \dots x_N^{t_1}), \dots, (a^{t_b}, x_1^{t_b}, x_i^{t_b} \dots x_N^{t_b})$, and repeat for the marginal distribution $p_A \otimes p_{\mathcal{X}} \sim (a^{t'_1}, x_1^{t_1}, x_i^{t_1} \dots x_N^{t_1}), \dots, (a^{t'_b}, x_1^{t_b}, x_i^{t_b} \dots x_N^{t_b})$, where $t'_i \neq t_i$.
4: Draw mini batch samples of length b from the joint distribution of the actions and the state with variable X_i missing. $p_{A,\mathcal{X}_{\setminus X_i}} \sim (a^{t_1}, x_1^{t_1} \dots x_N^{t_1}), \dots, (a^{t_b}, x_1^{t_b} \dots x_N^{t_b})$, and repeat for the marginal distribution $p_A \otimes p_{\mathcal{X}_{\setminus X_i}} \sim (a^{t'_1}, x_1^{t_1}, \dots x_N^{t_1}), \dots, (a^{t'_b}, x_1^{t_b} \dots x_N^{t_b})$, where $t'_i \neq t_i$.
5: $I(A; \mathcal{X}) = \frac{1}{b}\sum_{j=1}^{b} F_\theta(a^{t_j}, x_1^{t_j}, x_i^{t_j} \dots x_N^{t_j})) - \frac{1}{b}\sum_{j=1}^{b} \log e^{F_\theta((a_1^{t'_j}, x_1^{t_j}, x_i^{t_j} \dots x_N^{t_j})}$
6: $I(A; \mathcal{X}_{\setminus X_i}) = \frac{1}{b}\sum_{j=1}^{b} F_{\theta_{\setminus X_i}}((a^{t_j}, x_1^{t_j} \dots x_N^{t_j})) - \frac{1}{b}\sum_{j=1}^{b} \log e^{F_{\theta_{\setminus X_i}}(a_1^{t'_j}, x_1^{t_j} \dots x_N^{t_j})}$
7: **end for**
8: **return** $I(A; \mathcal{X}) - I(A; \mathcal{X}_{\setminus X_i})$

D Proof of Non-Negativity of $\Phi_{X_i\mathcal{X}\to A}$

In this section, we prove the non-negativity of $\Phi_{X_i;\mathcal{X}\to A}$. We first write the full expression of the measure as follows:

$$\Phi_{X_i;\mathcal{X}\to A} = -\int_{A\times\mathcal{X}} p_{A,\mathcal{X}}(a^t, x_1^t, x_i^t \dots x_N^t) \log \frac{p_{A,\mathcal{X}_{\setminus X_i}}(a^t, x_1^t \dots x_N^t), p_{\mathcal{X}}(x_1^t, x_i^t \dots x_N^t)}{p_{A,\mathcal{X}}(a^t, x_1^t, x_i^t \dots x_N^t), p_{\mathcal{X}_{\setminus X_i}}(x_1^t \dots x_N^t)} dA \times \mathcal{X}. \tag{8}$$

Figure 11: Empirical verification of Condition 1. For each environment, we plot the redundancy $-R(\mathcal{P}_1; \mathcal{P}_2 \to A)$ for the 10 most redundant subset pairs. The dashed line indicates $H(A)$; bars entering the shaded region would indicate a violation of Condition 1. In accordance with many redundancy synergy indexes, negative values indicate synergy positive values indicate redundancy. No violations are observed across any environment.

By applying Jensen's inequality we can then write:

$$
\begin{aligned}
\Phi_{X_i;\mathcal{X}\to A} &\geq -\log \int_{A\times\mathcal{X}} p_{A,\mathcal{X}}(a^t, x_1^t, x_i^t \dots x_N^t) \frac{p_{A,\mathcal{X}_{\setminus X_i}}(a^t, x_1^t \dots x_N^t), p_{\mathcal{X}}(x_1^t, x_i^t \dots x_N^t)}{p_{A,\mathcal{X}}(a^t, x_1^t, x_i^t \dots x_N^t), p_{\mathcal{X}_{\setminus X_i}}(x_1^t \dots x_N^t)} dA \times \mathcal{X} \\
&\geq -\log \int_{A\times\mathcal{X}} \frac{p_{A,\mathcal{X}_{\setminus X_i}}(a^t, x_1^t \dots x_N^t), p_{\mathcal{X}}(x_1^t, x_i^t \dots x_N^t)}{p_{\mathcal{X}_{\setminus X_i}}(x_1^t \dots x_N^t)} dA \times \mathcal{X} \\
&\geq -\log(1) \\
&\geq 0,
\end{aligned}
\tag{9}
$$

therefore proving the non-negativity of $\Phi_{X_i;\mathcal{X}\to A}$. □

E Proof of Lemma 1

The steps of this proof can be summarized as follows. Firstly, we show that a single element of a subset for which we observe CPMCR with another variable or subset of variables will not be included in the set $\mathcal{X}_\Phi$. Then, we will demonstrate that this applies to all the elements of both subsets. Finally, we use the definition of CPMCR to prove Lemma 1.

Let there be a case of CPMCR between two subsets $\mathcal{P}, \mathcal{P}' \subseteq \mathcal{X}$ such that:

$$\begin{aligned} H(A|\mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P} \cup \mathcal{P}') &= H(A|\mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}) \\ &= H(A|\mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}') \\ &< H(A|\mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')}). \end{aligned} \tag{10}$$

Let $\mathcal{P}_{\setminus\mathcal{P}} \in \mathscr{P}(\mathcal{X}_{\setminus\mathcal{P}})$ and $\mathcal{P}_{\setminus\mathcal{P}'} \in \mathscr{P}(\mathcal{X}_{\setminus\mathcal{P}'})$. Given Equation 2 and Equation 10, it must be true that:

$$H(A|\mathcal{P}_{\setminus\mathcal{P}'} \cup \mathcal{P}') \leq H(A|\mathcal{P}_{\setminus\mathcal{P}'} \cup \{P'\}), \tag{11}$$

where $P' \in \mathcal{P}'$. Combining Equations 10 and 11 leads us too:

$$H(A|\mathcal{P}_{\setminus\mathcal{P}} \cup \mathcal{P}) \leq H(A|\mathcal{P}_{\setminus\mathcal{P}} \cup \{P'\}), \tag{12}$$

we now substitute in $\mathcal{X}_{\setminus(\mathcal{P},P')} = \mathcal{P}_{\setminus\mathcal{P}}$, leading to:

$$\begin{aligned} H(A|\mathcal{X}_{\setminus(\mathcal{P},P')} \cup \mathcal{P}) &\leq H(A|\mathcal{X}_{\setminus(\mathcal{P},P')} \cup \{P'\}) \\ H(A|\mathcal{X}_{\setminus P'}) &\leq H(A|\mathcal{X}_{\setminus\mathcal{P}}) \\ &\leq H(A|\mathcal{X}), \end{aligned} \tag{13}$$

and therefore:

$$H(A|\mathcal{X}_{\setminus P'}) - H(A|\mathcal{X}) \leq 0. \tag{14}$$

Due to the non-negativity proof presented in Section D, we obtain $\Phi_{P';\mathcal{X}\to A} = 0$, and consequently element P' is not included in the set $\mathcal{X}_\Phi$. This holds $\forall P \in \mathcal{P}$ and $\forall P' \in \mathcal{P}'$. Consequently, neither of the elements of the subset $\mathcal{P}$ or of the subset $\mathcal{P}'$ will be included in the set $\mathcal{X}_\Phi$. Therefore, $\mathcal{X}_\Phi = \mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')}$, but, through Equation 10, we obtain $H(A|\mathcal{X}_\Phi) > H(A|\mathcal{X})$. □

F Proof of Theorem 1

We outline this proof in three steps, which we present as lemmas. The first of these lemmas characterizes the variables that are not to be included in $\mathcal{X}_\Phi$. By means of the other two, we demonstrate that by not adding these variables to the set $\mathcal{X}_\Phi$, we still satisfy $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$.

Lemma 2. *Let us assume that there exists variables X_i in $\mathcal{X}$ that transfers no entropy to the actions A (satisfying $\Phi_{X_i;\mathcal{X}\to A} = 0$).*

- *These variables satisfy one of either:*

$$H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_i}) < H(A|X_i) < H(A)), \tag{15}$$

  or

$$H(A|\{X_i\} \cup \mathcal{P}_{\setminus X_i}) = H(A|\mathcal{P}_{\setminus X_i})); \tag{16}$$

- *assuming no cases of CPMCR ($\neg\Psi(A|\mathcal{X})$), the following also holds:*

$$\begin{aligned} \Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi(A|\mathcal{X}) \Leftrightarrow & H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_i}) < H(A|X_i) < H(A) \quad \text{or} \\ & H(A|\{X_i\} \cup \mathcal{P}_{\setminus X_i}) = H(A|\mathcal{P}_{\setminus X_i}). \end{aligned} \tag{17}$$

Proof. Let us assume a case of CPMCR such that $\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})$, as this naturally leads to $\Phi_{X_i;\mathcal{X}\to A} = 0$. We now show that to satisfy the condition $\neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})$, the variable X_i must satisfy Equation 15 or Equation 16.

We first note that we can re-write $\Phi_{X_i;\mathcal{X}\to A} = 0$ as $H(A|\mathcal{X}_{\setminus X_i}) = H(A|\mathcal{X})$ and therefore the condition $\Phi_{X_i;\mathcal{X}\to A} = 0$ is already satisfied in the definition of CPMCR. Consequently, it is possible to write:

$$\begin{aligned}&\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})\\ \Leftrightarrow &(X_i \in \mathcal{X}, \exists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i}) : H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus \mathcal{P}'}) = H(A|\mathcal{X}_{\setminus X_i}) < H(A|\mathcal{X}_{\setminus (X_i,\mathcal{P}')}) \quad \& \\ &\psi_{\mathcal{P}'}).\end{aligned} \tag{18}$$

Suppose that we violate the condition $\psi'_{\mathcal{P}}$ by adding a non-informative variable V_{rand} to the set $\mathcal{P}'$. Although, this would lead to $\neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})$, it would still be true that $\Psi_{X_i,\mathcal{P}'_{\setminus V_{rand}}}(A|\mathcal{X})$. Consequently, there will always exist $\mathcal{P}' \in \mathscr{P}$ such that $\psi'_{\mathcal{P}}$ is satisfied. Given there always exists a scenario in which $\psi'_{\mathcal{P}}$ is true, it cannot be used to induce the condition $\neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})$. Consequently, for clarity, we remove it from the presentation and we write:

$$\begin{aligned}&\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})\\ &\Leftrightarrow (X_i \in \mathcal{X}, \exists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i}) : H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus \mathcal{P}'}) = H(A|\mathcal{X}_{\setminus X_i}) < H(A|\mathcal{X}_{\setminus (X_i,\mathcal{P}')})).\end{aligned} \tag{19}$$

The next step in this proof is to rewrite Equation 19 so that it satisfies $\neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})$. This will be the case if the equality that relates $H(A|\mathcal{X}_{\setminus \mathcal{P}'}) = H(A|\mathcal{X}_{\setminus X_i})$ or $H(A|\mathcal{X}_{\setminus X_i}) < H(A|\mathcal{X}_{\setminus (X_i,\mathcal{P}')})$ no longer holds. These are the only ways to derive $\neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})$ without leading to $\neg\Phi_{X_i;\mathcal{X}\to A}$.

We now assume the equality that relates $H(A|\mathcal{X}_{\setminus \mathcal{P}'}) = H(A|\mathcal{X}_{\setminus X_i})$ no longer holds. We obtain a variable X_i and subset $\mathcal{P}'$ that do not provide the same information about the target variable $H(A|\mathcal{X}_{\setminus \mathcal{P}'}) > H(A|\mathcal{X}_{\setminus X_i})$. Substituting this into Equation 19 we have:

$$\begin{aligned}&\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})\\ &\Leftarrow (X_i \in \mathcal{X}, \exists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i}) : H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_i}) < H(A|\mathcal{X}_{\setminus \mathcal{P}'}) < H(A|\mathcal{X}_{\setminus (X_i,\mathcal{P}')})),\end{aligned} \tag{20}$$

Suppose we rewrite Equation 20 letting $\mathcal{P}' = \mathcal{X}_{\setminus X_i}$. We obtain:

$$\begin{aligned}&\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})\\ &\Leftarrow (H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_i}) < H(A|\mathcal{X}_{\setminus \mathcal{X}_{\setminus X_i}}) \leq H(A|\mathcal{X}_{\setminus \{X_i,\mathcal{X}_{\setminus X_i}\}}))\\ &\Leftarrow (H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_i}) < H(A|X_i) < H(A)).\end{aligned} \tag{21}$$

Hence, the variables that satisfy Equation 16 also satisfy $\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})$. Now, let us suppose that the inequality $H(A|\mathcal{X}_{\setminus X_i}) < H(A|\mathcal{X}_{\setminus (X_i,\mathcal{P}')})$ no longer holds. Specifically, we assume neither the subset $\mathcal{P}'$ or variable X_i provide any information about the target variable. In this case it is true that $H(A|\mathcal{X}_{\setminus X_i}) = H(A|\mathcal{X}_{\setminus (X_i,\mathcal{P}')})$ and we can write the following:

$$\begin{aligned}&\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})\\ &\Leftarrow (X_i \in \mathcal{X}, \exists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i}) : H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_i}) = H(A|\mathcal{X}_{\setminus \mathcal{P}'}) = H(A|\mathcal{X}_{\setminus (X_i,\mathcal{P}')})).\end{aligned} \tag{22}$$

We now substitute in $\mathcal{P}_{\setminus (X_i,\mathcal{P}')} = \mathcal{X}_{\setminus (X_i,\mathcal{P}')}$ in Equation 22 such that:

$$\begin{aligned}&\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X})\\ &\Leftarrow (X_i \in \mathcal{X}, \exists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i}) : H(A|\mathcal{X}) = H(A|\mathcal{P}_{\setminus (X_i,\mathcal{P}')} \cup \mathcal{P}') = H(A|\mathcal{P}_{\setminus (X_i,\mathcal{P}')} \cup X_i) = H(A|\mathcal{P}_{\setminus (X_i,\mathcal{P}')}).\end{aligned} \tag{23}$$

Given $\mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i})$ and $\mathcal{P}_{\setminus (X_i,\mathcal{P}')} \in \mathscr{P}(\mathcal{X}_{\setminus (X_i,\mathcal{P}')})$, it must be so that $\mathcal{P}_{\setminus (X_i,\mathcal{P}')} \cup \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i})$. Consequently, we replace $\mathcal{P}_{\setminus (X_i,\mathcal{P}')} \cup \mathcal{P}'$ with $\mathcal{P}_{\setminus X_i}$ in Equation 23 to give

$$\begin{aligned}\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X}) \Leftarrow (X_i \in \mathcal{X}, \exists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{\setminus X_i}) :\\ H(A|\mathcal{P}_{\setminus X_i} \cup X_i) = H(A|\mathcal{P}_{\setminus X_i}) = H(A|\mathcal{P}_{\setminus (X_i,\mathcal{P}')} \cup X_i) = H(A|\mathcal{P}_{\setminus (X_i,\mathcal{P}')})).\end{aligned} \tag{24}$$

Since $\mathcal{P}_{\setminus(X_i,\mathcal{P}')} \subseteq \mathcal{P}_{\setminus X_i}$, we do not include the last two equivalences as they are sub-conditions of the first equivalence. Therefore, we can write

$$\Phi_{X_i;\mathcal{X}\to A} = 0 \quad \& \quad \neg\Psi_{X_i,\mathcal{P}'}(A|\mathcal{X}) \Leftarrow H(A|\mathcal{P}_{\setminus X_i} \cup X_i) = H(A|\mathcal{P}_{\setminus X_i})). \tag{25}$$

Combining the statements in Equations 25 and 21 completes the proof of Lemma 2. □

According to Lemma 2, variables that satisfy Equation 16 are not included in the set $\mathcal{X}_\Phi$, despite this, we now prove that $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$ still holds.

Lemma 3. *Let us assume there exists a non-empty subset of variables in $\mathcal{X}$ that satisfies Equation 15. The following holds: $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$.*

Proof. To prove Lemma 3, we let variables X_j and X_k satisfy the property expressed by Equation 16, and show that despite removing them from $\mathcal{X}$ to form $\mathcal{X}_\Phi$, $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$ is still satisfied. We then demonstrate that this result generalizes to cases with more than two variables.

In accordance with Lemma 2, if we remove X_j from $\mathcal{X}$ we have $H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_j})$. Suppose that we now repeat this process, but instead, we remove X_k from $\mathcal{X}_{\setminus X_j}$, as this will obtain $\mathcal{X}_\Phi$. To begin, we let $\mathcal{P}_{\setminus X_k} \in \mathscr{P}(\mathcal{X}_{\setminus X_k})$, such that:

$$H(A|\mathcal{P}_{\setminus X_k}) = H(A|\{X_k\} \cup \mathcal{P}_{\setminus X_k}). \tag{26}$$

We now substitute it in $\mathcal{P}_{\setminus X_k} = \mathcal{X}_{\setminus\{X_k,X_j\}}$:

$$H(A|\mathcal{X}_{\setminus\{X_k,X_j\}}) = H(A|\{X_k\} \cup \mathcal{X}_{\setminus\{X_k,X_j\}}), \tag{27}$$

because $\{X_k\} \cup \mathcal{X}_{\setminus\{X_k,X_j\}} = \mathcal{X}_{\setminus X_j}$ we have:

$$\begin{aligned} H(A|\mathcal{X}_{\setminus\{X_k,X_j\}}) &= H(A|\mathcal{X}_{\setminus X_j}) \\ &= H(A|\mathcal{X}). \end{aligned} \tag{28}$$

However, given that only the variables X_j and X_k satisfy the property expressed by Equation 16, we have $\mathcal{X}_{\setminus\{X_k,X_j\}} = \mathcal{X}_\Phi$ and therefore $H(A|\mathcal{X}_\Phi) = H(A|\mathcal{X})$. Consequently, we have proved Lemma 3 in the two variable case.

We now generalize this proof beyond the two variable case. Let variable X_l also satisfy the property outlined in Equation 16, we now have $\mathcal{X}_{\setminus\{X_k,X_j,X_l\}} = \mathcal{X}_\Phi$. We complete this proof, by first using the property outlined in Equation 16 to show $H(A|\mathcal{X}_{\setminus\{X_k,X_j\}}) = H(A|\mathcal{X}_{\setminus\{X_k,X_j,X_l\}})$, before then combining this with the proof for the two variable case to show $H(A|\mathcal{X}_\Phi) = H(A|\mathcal{X})$. We can repeatedly apply this proof up to the N variable case. □

According to Lemma 2, the variables that satisfy Equation 15 are not included in the set $\mathcal{X}_\Phi$. Despite this, provided no cases of CPMCR exist in the set $\mathcal{X}$, it is possible to prove that $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$ still holds. More formally, we demonstrate the following lemma:

Lemma 4. *Let us assume there exists a non-empty subset of variables in $\mathcal{X}$ that satisfies Equation 15. Provided no cases of CPMCR exist in $\mathcal{X}$, the following holds: $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$.*

Proof. We let variables X_j and X_k satisfy the property outlined in Equation 15, and show that if removing them to form $\mathcal{X}_\Phi$ leads to $H(A|\mathcal{X}) < H(A|\mathcal{X}_\Phi)$, a contradiction arises. Therefore, it must be true that $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$. Subsequently, we generalize to the N variable case.

If Equation 15 holds for variables X_j and X_k, so does $H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_j}) = H(A|\mathcal{X}_{\setminus X_k})$. Now we assume that Lemma 4 is not true. In this case, we have: $H(A|\mathcal{X}) < H(A|\mathcal{X}_\Phi)$. Currently, we are only interested in the two-variable case, so this becomes $H(A|\mathcal{X}) < H(A|\mathcal{X}_{\setminus\{X_j,X_k\}})$. Consequently, we have:

$$H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_j}) = H(A|\mathcal{X}_{\setminus X_k}) < H(A|\mathcal{X}_{\setminus\{X_j,X_k\}}). \tag{29}$$

Equation 29 reveals that in order to verify the condition $H(A|\mathcal{X}) < H(A|\mathcal{X}_\Phi)$, there must be CPMCR, such that $\Psi_{X_j,X_k}(A|\mathcal{X})$ (ψ here is invalid as we are concerned with a single variable). Consequently, a contradiction has arisen; therefore, we have proven Lemma 4 in the two-variable case.

We now generalize this result to the N variable case. Let there be a third variable, X_l, that also satisfies the property expressed by Equation 15. In this case, it must be true that: $H(A|\mathcal{X}) = H(A|\mathcal{X}_{\setminus X_j}) = H(A|\mathcal{X}_{\setminus X_k}) = H(A|\mathcal{X}_{\setminus X_l})$. Now we reapply an identical method to demonstrate that, if $H(A|\mathcal{X}) < H(A|\mathcal{X}_\Phi) = H(A|\mathcal{X}_{\setminus\{X_j,X_k,X_l\}})$, we must have a case of CPMCR, such that $\Psi_{X_j,X_k,X_l}(A|\mathcal{X})$. Again, a contradiction has arisen, and this proves the lemma in the case of N variables. □

Theorem 1. *Please refer to Section 6.4.*

Proof. In Lemma 2, we have shown that variables that provide no information about the actions (satisfying Equation 16) and variables that provide redundant information about the actions (satisfying Equation 15) will not be included in the set $\mathcal{X}_\Phi$. It is therefore trivial to see that the remaining variables in the set $\mathcal{X}_\Phi$ provide non-redundant information about the actions, and, therefore, satisfy $H(A|\mathcal{X}_{\Phi\setminus X_i}) > H(A|\mathcal{X}_\Phi)$. This means that it is not possible to remove any more variables from $\mathcal{X}_\Phi$ without increasing the value of $H(A|\mathcal{X}_\Phi)$. Consequently, the set $\mathcal{X}_\Phi$ satisfies the condition $|\mathcal{X}_\Phi| = \min_{H(A|\mathcal{P})=H(A|\mathcal{X}_\Phi)} |\mathcal{P}|$, where $\mathcal{P} \in \mathscr{P}(\mathcal{X})$. Combining this information with the results of Lemmas 4 and 5, in which we have shown that despite removing these variables the expression $H(A|\mathcal{X}) = H(A|\mathcal{X}_\Phi)$ still holds, completes the proof of Theorem 1. □

G Proof of Theorem 2

To carry out this proof we first present Lemma 6, which shows that Algorithm 1 effectively deals with cases of CPMCR between subsets of equal cardinality. Using this result alongside Theorem 1, we demonstrate that the proof is valid under Condition 1.

Lemma 6. *Let $\mathcal{P} \in \mathscr{P}(\mathcal{X}_{A_1})$, where $\mathcal{X}_{A_1}$ is generated using Algorithm 1. The following holds*[7]*:*

$$\forall \mathcal{P} \in \mathscr{P}(\mathcal{X}_{A_1}) \nexists \mathcal{P}' \in \mathscr{P}(\mathcal{X}_{A_1}) : \Psi(A|\mathcal{X}) \quad \& \quad |\mathcal{P}| = |\mathcal{P}'| \quad \& \quad \mathcal{P} \neq \mathcal{P}'. \tag{30}$$

Proof. We prove Lemma 6 by demonstrating that Algorithm 1 is able to resolve cases of CPMCR between subsets with equal cardinality. We consider two subsets $\mathcal{P}$ and $\mathcal{P}'$ with equal cardinality ($|\mathcal{P}| = |\mathcal{P}'|$) and CPMCR between them ($\Psi_{\mathcal{P},\mathcal{P}'}(A|\mathcal{X})$). We can write the set $\mathcal{X}$ that contains these subsets as: $\mathcal{X} = \{P_1, P_1', P_2, P_2' \dots P_N, P_N'\} \cup \mathcal{X}_{\setminus(\mathcal{P},\mathcal{P}')}$, where $P_i \in \mathcal{P}$ and $P_i' \in \mathcal{P}'$. We can now apply Algorithm 1: we begin by calculating $\Phi_{P_1;\mathcal{X}\to A} = 0$, in accordance with Lemma 1. Consequently, P_1 gets removed from the set $\mathcal{X}$, so we can write $\mathcal{X} = \mathcal{X}^0_{\setminus P_1}$, where here $\mathcal{X}^0$ indicates $\mathcal{X}$ prior to any algorithmic operations. We now calculate $\Phi_{P_1';\mathcal{X}\to A}$, as we iterate through the set $\mathcal{X}$ ($\mathcal{X}^0_{\setminus P_1}$). To understand the result of this calculation, we first point out that the subsets $\mathcal{P}$ and $\mathcal{P}'$ are defined such that their respective elements combine to provide identical information about the actions (see Equation 2). Therefore, for a single element, the following holds:

$$H(A|\mathcal{X}^0_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}') < H(A|\mathcal{X}^0_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}'_{\setminus P_i'}). \tag{31}$$

Consequently, the following must also be true:

$$\begin{aligned} H(A|\mathcal{X}^0_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}' \cup \mathcal{P}_{\setminus P_1}) &< H(A|\mathcal{X}^0_{\setminus(\mathcal{P},\mathcal{P}')} \cup \mathcal{P}'_{\setminus P_1'} \cup \mathcal{P}_{\setminus P_1}) \\ H(A|\mathcal{X}^0_{\setminus P_1}) &< H(A|\mathcal{X}^0_{\setminus\{P_1,P_1'\}}) \\ H(A|\mathcal{X}) &< H(A|\mathcal{X}_{\setminus P_1'}). \end{aligned} \tag{32}$$

Thus, $\Phi_{P_1';\mathcal{X}\to A} > 0$. We now include P_1' in $\mathcal{X}_{A_1}$ and we do not remove it from $\mathcal{X}$. We will then consider variable P_2, and exclude it from $\mathcal{X}$ for similar reasons to P_1. The variable P_2' will then remain in the set $\mathcal{X}$ and be added to $\mathcal{X}_{A_1}$ for reasons similar to P_1'. This process will be repeated until we have removed all the elements of the subset $\mathcal{P}$ from the set $\mathcal{X}$, while the elements of the set $\mathcal{P}'$ have been added to $\mathcal{X}_{A_1}$.

[7] In this lemma, we adopt the definition of $\Psi(A|\mathcal{X})$ provided in Equation 3.

We have therefore proved that Algorithm 1 reliably addresses cases of CPMCR between subsets with equal cardinality. □

Theorem 2. *See Section 6.4.*

Proof. Given Condition 1 and the proof of Lemma 6 outlined in Appendix G, we can disregard cases of CPMCR, and, therefore, directly apply Theorem 1 to complete this proof. □

H Tit-For-N-Tats Extra Graphs

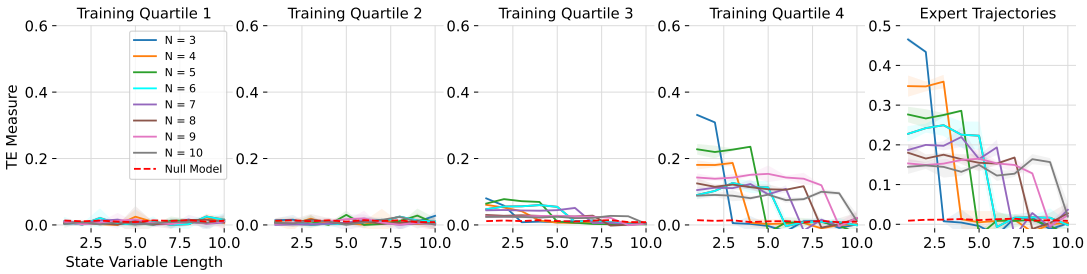

(A)

Figure 12: In subfigure (A), we use $\Phi_{X_i;\mathcal{X}\to A}$ to investigate how entropy is transferred from observable state variables to actions during different stages of training.

I Implementation Details

I.1 Synthetic Experiments

Transfer entropy measure estimation. For the estimation of the amount of entropy that each variable transfers to the actions, we use a standard feed-forward Neural Network with a ReLu activation function and with one hidden layer with 50 nodes. We select the following values for the parameters in Algorithm 2: $b = 10000$, $N = 20000$, with a learning rate of 0.01.

Baselines. For the calculation of UMFI, we select the exact values of the hyperparameters as those used in (Janssen et al., 2023). As far as PI is concerned, we adopt the following values: $\alpha = 0.01$, and 1000 runs. For the calculation of PI itself, we employed Ridge (Pedregosa et al., 2011), as implemented in Scickit-Learn. In fact, in this case, hyperparameter tuning has no impact on the final UMFI and PI values.

I.2 The Secret Key Game

Trajectory generation. To carry out the polynomial interpolation in the Secret Key Game, we use the NumPy's `polynomials` package `https://numpy.org/doc/stable/reference/routines.polynomials.package.html`. An expert trajectory for this game is one in which the agent achieves a score of zero. For handling the state spaces of lengths 25 and 50, we employ one-step temporal difference Actor-Critic learning during the execution of the game.

The Actor network is updated using Equation $\theta'_{actor} \leftarrow \theta_{actor} + \alpha_{actor}\gamma^t A \nabla \ln \pi(a^t|s^t, \theta_{actor})$, where $\delta = r + \gamma v(s^t, \theta_{critic}) - v(s^{t+1}, \theta_{critic})$ (Sutton & Barto, 2018). The Critic network is updated using the following equation: $\theta'_{critic} \leftarrow \theta_{critic} + \alpha_{critic} MSE(\delta)$. Both networks are characterized by a single fully connected layer of size 64 with a ReLu activation function. The other values of the hyperparameters of the networks are $\gamma = 0.99, \alpha_{actor} = 0.0001, \alpha_{critic} = 0.001$.

Transfer entropy measure estimation. For the estimation of the amount of entropy that each of the keys transfers to the actions, we again use a network with one hidden layer with 50 nodes and a ReLu activation function. For 25 keys, we use $b = 5000$, $N = 2000$ in Algorithm 2. Instead, for a state with

50 keys, we require the following parameters for Algorithm 2 $b = 10000$, $N = 20000$. For both, we used a learning rate of 0.01.

Baselines. For both 25 and 50 keys, we use the exact hyperparameter values specified in (Janssen et al., 2023) for calculating UMFI, as they are sufficient to identify the correct keys in the game. For the PI calculation, we employ the same hyperparameter values as those used in the synthetic data experiments.

Evaluation of designed states. We employ a one-step temporal difference Actor-Critic architecture when evaluating each state. Both networks have a single fully connected layer of size 64 and uses a ReLu activation function. We employ the following values for the parameters: $\gamma = 0.99, \alpha_{actor} = 0.0001, \alpha_{critic} = 0.001$, considering 20,000 episodes of the game (in this game each episode requires only one iteration).

I.3 OpenAI's Gym: Cart Pole

Trajectory generation. We play the game until convergence. An expert trajectory is defined as one that achieves a cumulative reward greater than 475, in line with the criteria proposed by the authors of the environments (Brockman et al., 2016). This is done using the same Actor-Critic architecture as described for the Secret Key Game. We introduce random variables into the state, each following a uniform probability distribution within the range $V_{rand_i} \in [-5, 5]$. However, if the variables are truly random, the specific range should be inconsequential.

Transfer entropy measure estimation. We use the same architecture as previously stated for the Secret Key Game, except in this case, we use the following parameters for Algorithm 2 $b = 10000$, $N = 4000$, and $\alpha = 0.01$.

Baselines. For these experiments when calculating UMFI we use the same values of hyperparameters as in (Janssen et al., 2023), except for the fact that we adopt 50 trees instead of 100. This was because we saw improved performance using 50 trees for the task at hand. For the PI calculation, we adopted the same values of the hyperparameters as those used for the synthetic data experiments.

Evaluation of designed states. When plotting the state performance curves shown in Figure 5, we use the same architecture as for the generation of the trajectories.

I.4 OpenAI's Gym: Lunar Lander

Trajectory generation. In accordance with Open AI's documentation, we consider expert trajectories as those which achieve a cumulative reward above 200. Again, the method is applied after playing the game until convergence. This is done using the same Actor-Critic architecture as described for the Secret Key Game. We inject three random variables into the state, in the same manner as described for the cart-pole game.

Transfer entropy measure estimation. We use the same architecture as previously stated, but here we use the following parameters for Algorithm 2: $b = 25000$, $N = 100000$, and a learning rate of 0.01.

Baselines. Again, for these experiments when calculating UMFI, we adopt the same values of the hyperparameters as used in (Janssen et al., 2023), except for the fact that we use 50 trees instead of 100 (since 50 trees are sufficient to separate the informative variables from the non-informative ones). For the PI calculation, we again adopted the same values of the hyperparameters as used for the synthetic data experiments.

Evaluation of designed states. For the evaluation of the training times corresponding to states composed of different numbers of variables, we use the same architecture as that used for generating trajectories. In this case, we only run the experiment for 3000 episodes.

I.5 OpenAI's Gym: Pendulum

Trajectory generation. Again we play the game until convergence. This is done using a PPO architecture, with the following Actor update rule:

$$\theta' \leftarrow \theta + \alpha\gamma^t \nabla min(\frac{\pi(a^t|s^t,\theta^t)}{\pi(a^t|s^t,\theta_k^t)}A^t, clip(\frac{\pi(a^t|s^t,\theta^t)}{\pi(a^t|s^t,\theta_k^t)}, 1-\epsilon, 1+\epsilon)A^t, \tag{33}$$

where the surrogate actor had its weights updated as described except it is multiplied by 0.95. For this algorithm, we use the following values for the hyperparameters: $\gamma = 0.99, \alpha = 0.0003, \epsilon = 0.2$. We also assume a mini-batch size of 64 and update our networks every 2048 steps. We use 10 epochs and an entropy coefficient of 0.001. The network has one hidden layer with 64 nodes and uses continuous hyperbolic tangent activation functions according to Schulman et al. (2017). We consider trajectories with cumulative rewards greater than -10 as expert trajectories.

Transfer entropy measure estimation. We use the same architecture to estimate the TE values, combined with the following parameters for Algorithm 2 $b = 5000$, $N = 4000$, with a learning rate of 0.01.

Baselines. For these experiments, when calculating UMFI we employ the same values of the hyperparameters as those used in (Janssen et al., 2023). For the PI calculation, we adopted the same hyperparameters as used for the synthetic data experiments.

Evaluation of designed states. We use the same PPO architecture when investigating the performance of different states as that used for generating the trajectories.

I.6 TFNT in the Iterated Prisoner's Dilemma

		Player 1	
		Cooperate	Defect
Player 2	Cooperate	2, 2	3, 0
	Defect	0, 3	1, 1

Figure 13: Iterated Prisoner's Dilemma payoff matrix.

Trajectory generation. We employ tabular Q-learning where $\alpha = 0.9$, and the value of ϵ decays from 1 to 0 during $40,000$ iterations ($\epsilon^t = \epsilon^{t-1} - (1/40,000)$), while γ is equal to 0.99. We play the game with the above payoff matrix until convergence. We use one-shot encoding in our implementation to generate the state at a certain time step. For example, let us assume that the last 9 action pairs could be represented by $s^t = [(C,D),(C,D),(C,C),(D,D),(C,D),(D,D),(C,D),(D,C),(D,C)]$, this can be re-written as: $s^t = [(2),(2),(0),(3),(2),(3),(2),(1),(1)]$.

Transfer entropy measure estimation. We use the same architecture as for the other experiments combined with the following values of the parameters for Algorithm 2: $b = 1000$, $N = 10000$, with a learning rate of 0.01.

Baselines. For these experiments, for the calculation of UMFI, we use the exact values of the hyperparameters as those in (Janssen et al., 2023). We select these values since hyperparameter tuning has no impact on the final UMFI values. For the PI calculation, we again adopt the same hyperparameters as those used for the synthetic data experiments.

Evaluation of designed states. We employ the same Q-learning architecture when evaluating the training performance of different state representations.

I.7 Null Model

The null model consists of a series of random integers between 0 and 1 that are added to the trajectories after their generation. For all experiments, the values of the hyperparameters used in Algorithm 2 to estimate $\Phi_{NM;\mathcal{X}\to A}$ are the same as for the variables of interest. For a variable to be considered as influencing the actions in a statistically significant way they must have a 95% chance of falling outside the range of the null model, and, therefore, $\mu_{Xi} - \frac{2\sigma_{Xi}}{\sqrt{10}} > \mu_{NM} + \frac{2\sigma_{NM}}{\sqrt{10}}$, for all variables included in $\mathcal{X}_*$.