Approximate Sufficient Statistics for Team Decision Problems

Alex Lemon
Department of Electrical Engineering
Stanford University
Stanford, CA 94035, USA
adlemon@stanford.edu

Sanjay Lall
Department of Aeronautics and Astronautics
Stanford University
Stanford, CA 94035, USA
lall@stanford.edu

Abstract

Team decision problems are one of the fundamental problems in decentralized decision making. Because team decision problems are NP-hard, it is important to find methods for reducing their complexity. Wu and Lall gave a definition of sufficient statistics for team decision problems, and demonstrated that these statistics are sufficient for optimality, and possess other desirable properties such as being readily updated when additional information becomes available. More recently, Lemon and Lall defined weak sufficient statistics for team decision problems, and showed that these statistics are sufficient for optimality and necessary for simultaneous optimality with respect to all cost functions. This prior work studied the extent to which the complexity of team decision problems can be reduced while maintaining exact optimality. However, when faced with a computationally difficult problem, we are often willing to sacrifice exact optimality for significant reductions in complexity. In this paper we define approximate sufficient statistics, which are a generalization of weak team sufficient statistics. Then we prove that these statistics are quantifiably close to being optimal.

Introduction

This paper is about multi-agent decentralized decision making. In particular, we consider team decision problems (TDPs), which are a static simplification of decentralized Markov decision problems (Dec-MDPs). While Markov decision problems (MDPs) are P-complete (Papadimitriou and Tsitsiklis 1987), and are therefore considered "easy" problems, decentralized Markov decision processes are NEXP-complete (Bernstein et al. 2002), and considered "hard" problems. Nonetheless, certain classes of Dec-MDPs and related problems are known to have explicit solutions (Kube and Zhang 1997; Hansen, Bernstein, and Zilberstein 2004; Wu and Lall 2010a; 2010b). The insight needed to construct these explicit solutions often arises from the consideration of sufficient statistics (Wu and Lall 2015; Wu 2013; Lessard and Lall 2015). Moreover, sufficient statistics have also been used to reduce the complexity of search algorithms for Dec-MDPs (Dibangoye et al. 2013; Oliehoek, Whiteson, and Spaan 2009; Oliehoek 2013). Intuitively, sufficient statistics are simplified or compressed versions of the measurements that allow us to make decisions that are just as good as decisions made using the unsimplified measurements. Thus, identifying sufficient statistics is often a crucial step in reducing the complexity of multi-agent problems. Although TDPs are a simplification of Dec-MDPs, solution techniques for TDPs sometimes provide approaches to MDPs, and, in particular, many of the ideas related to sufficient statistics for TDPs generalize to dynamic problems (Wu and Lall 2015).

Team decision problems model a scenario in which a team of agents interact with a stochastic system. Each agent receives a measurement, and must choose an action solely on the basis of its measurement. (In particular, the agents do not share their measurements.) The objective of the team is to minimize the expected cost, where the cost is a function of the state of the system with which the team interacts and the actions of the agents. This class of problems was introduced in the context of organizations (Marschak 1955), and is one of the fundamental building blocks of decentralized-control problems (Ho and Chu 1972). An early paper showed that optimal policies for team decision problems are linear in the special case when the random variables are Gaussian, and the cost functions are quadratic (Radner 1962).

However, team decision problems are known to be hard in the general case (Tsitsiklis and Athans 1985), so finding ways to reduce the complexity of these problems is essential, and approximate techniques may be acceptable if they give a sufficiently large reduction in complexity. As in the case of classical single-player decision problems, sufficient statistics for TDPs are meant to encapsulate all of the information in the measurements needed to make optimal decisions. If these statistics exclude extraneous information, then the task of finding an optimal policy for the decision problem may be considerably simplified. In particular, identifying sufficient statistics may simplify the search for optimal policies (Lessard and Lall 2015). A class of team sufficient statistics (Wu 2013; Wu and Lall 2014a; 2014b; 2015) is known to be sufficient for optimality, and possess many other desirable properties. A different class of weak sufficient statistics (Lemon and Lall 2015) has been shown to be not only sufficient for optimality, but also necessary for simultaneous optimality with respect to all cost functions.

In this paper we generalize the idea of weak team sufficient statistics to define approximate sufficient statistics for team decision problems. We then prove that decisions based on these approximate sufficient statistics are quantifiably close to being optimal. The paper concludes with numerical examples in which approximate sufficient statistics allow us to determine which of a given set of statistics is most valuable for making decisions.

Notation and Problem Formulation

Given a finite set S, a function $f : S \to \mathbf{R}$, and a scalar $\alpha \in [1, \infty]$, the α-norm of f is

$$\|f\|_\alpha = \left(\sum_{s \in S} |f(s)|^\alpha \right)^{1/\alpha}.$$

The dual norm of the α-norm is the β-norm, where $\beta \in [1, \infty]$ is the unique number satisfying $\alpha^{-1} + \beta^{-1} = 1$. If the β-norm is the dual norm of the α-norm, we will sometimes write $\|f\|_{\alpha *}$ for $\|f\|_\beta$.

We use $\operatorname{conv}(S)$ to denote the convex hull of a set S. Given a finite set Ω, let $\Delta(\Omega)$ be the set of all probability distributions on Ω:

$$\Delta(\Omega) = \left\{ p : \Omega \to \mathbf{R} \;\middle|\; \begin{array}{l} \sum_{\omega \in \Omega} p(\omega) = 1, \\ p(\omega) \geq 0 \text{ for all } \omega \in \Omega \end{array} \right\}.$$

The total-variation distance between $p, q \in \Delta(\Omega)$ is

$$\begin{aligned} \mathrm{d}_{\mathrm{TV}}(p, q) &= \max_{A \subset \Omega} |p(A) - q(A)| \\ &= \frac{1}{2} \sum_{\omega \in \Omega} |p(\omega) - q(\omega)| \\ &= \frac{1}{2} \|p - q\|_1. \end{aligned}$$

Suppose $A_1, \ldots, A_n$ and $B_1, \ldots, B_n$ are sets. Define $A = A_1 \times \cdots \times A_n$ and $B = B_1 \times \cdots \times B_n$. We say that a function $f : A \to B$ is diagonal if there exist functions $f_j : A_j \to B_j$ for $j = 1, \ldots, n$ such that

$$f(a) = (f_1(a_1), \ldots, f_n(a_n)).$$

(A diagonal function is called a decentralized function in some fields: for example, a diagonal control policy is often called a decentralized control policy in the control-theory literature, or a decentralized joint policy in the Dec-POMDP community.) The set of all diagonal functions from A to B is denoted

$$\mathcal{D}(A, B) = \{ f : A \to B \mid f \text{ is diagonal} \}.$$

Given sets X, Y, and Z, and functions $f : Y \to Z$ and $g : X \to Y$, the composition of f and g is written $f \circ g$. In particular, if $y : \Omega \to Y$ is a random variable on a sample space Ω, S is a set, and $h : Y \to S$ is a function, then $h \circ y : \Omega \to S$ is also a random variable on Ω.

Team Decision Problems and Sufficient Statistics

Suppose a team of n agents interacts with a stochastic system. Let the random variable $x : \Omega \to X$ represent the state of the system, where Ω is a finite sample space. For $j = 1, \ldots, n$, the jth agent receives two measurements, represented by the random variables $y_j : \Omega \to Y_j$ and $s_j : \Omega \to S_j$. (Later in the paper we will assume that s_j is a statistic of y_j; however, our earlier results hold more generally, so we do not start out with this assumption.) Agent j must use its measurement (either y_j or s_j) to choose an action $u_j \in U_j$. For convenience we define the sets $Y = Y_1 \times \cdots \times Y_n$, $S = S_1 \times \cdots \times S_n$, and $U = U_1 \times \cdots \times U_n$. Let $C : X \times U \to \mathbf{R}$ be a cost function. The objective of the team is to minimize the expected cost. A team decision problem in y is an optimization problem of the form

$$C^\star(y) = \min_{\phi \in \mathcal{D}(Y, U)} \mathbf{E}\, C(x, \phi \circ y).$$

where the cost is a function of the system state x, and the actions $\phi_1(y_1), \ldots, \phi_n(y_n)$. We also consider team decision problems in s:

$$C^\star(s) = \min_{\psi \in \mathcal{D}(S, U)} \mathbf{E}\, C(x, \psi(s)).$$

We are interested in comparing $C^\star(y)$ and $C^\star(s)$. Earlier work has focused on determining when it is possible to guarantee that $C^\star(s) \leq C^\star(y)$. In particular, team sufficient statistics are known to be sufficient for optimality, and possess a number of other desirable properties (Wu 2013; Wu and Lall 2014a; 2014b; 2015). Similarly, weak team sufficient statistics have been shown to be not only sufficient for optimality, but also necessary for simultaneous optimality with respect to all cost functions (Lemon and Lall 2015). (However, weak team sufficient statistics have not been as thoroughly studied, and it is unknown whether they possess some of the attractive properties of team sufficient statistics.)

The approximate sufficient statistics that we define in this paper are a natural generalization of weak team sufficient statistics. Therefore, we briefly review the existing theory of team sufficient statistics. Before presenting the formal definitions, we give a motivating example that illustrates some of the intuitive meaning of team sufficient statistics. This example shows how the information structure in a problem is often reflected in the sufficient statistics for the problem.

Example 1. *Suppose x_1, w_2, w_3, v_1, v_2, and v_3 are independent random variables. We assume that*

$$x_2 = f_{21}(x_1, w_2) \quad \text{and} \quad x_3 = f_{31}(x_1, w_3)$$

for some functions f_{21} and f_{31}. The information structure of this problem is shown in fig. 1. This network reflects the fact that x_1 affects x_2 and x_3, but not vice versa, and that x_2 and x_3 are conditionally independent given x_1. Suppose a team of three agents interacts with a system whose state is x. The team receives the measurements z_1, z_2, and z_3, where

$$z_j = h_j(x_j, v_j)$$

for some functions h_1, h_2, and h_3. In particular, the first agent receives the measurement $y_1 = z_1$, the second agent receives the measurement $y_2 = (z_1, z_2)$, and the third agent receives the measurement $y_3 = (z_1, z_3)$. It is possible to show (Wu 2013) that team sufficient statistics for x given $y = (y_1, y_2, y_3)$ are $(s_1, (s_1, s_2), (s_1, s_3))$, where s_3 is

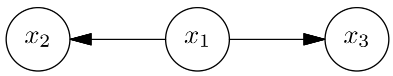

Figure 1: an information structure

(classically) sufficient for (x_1, x_3) given y_3, s_2 is (classically) sufficient for (x_1, x_2) given y_2, and s_1 is (classically) sufficient for s_2 given y_1, and for s_3 given y_1. The proof of this result relies on a calculus of team sufficient statistics, which allows us to construct team sufficient statistics from classically sufficient statistics.

Recall that one characterization of sufficient statistics for classical single-agent decision problems states that if x and y are jointly distributed random variables, then a statistic s of y is a sufficient statistic for x given y if x and y are conditionally independent given s. Thus, the first step in defining sufficient statistics for team decision problems is generalizing the notion of conditional independence to the multi-agent case. Wu and Lall gave the following definition of team independence.

Definition 2. *We say that y is team independent of x given s, and write $y \perp\!\!\!\perp x \mid s$, if*

$$p(x, y, s) \in \operatorname{conv}\{p(x, g \circ s, s) \mid g \in \mathcal{D}(S, Y)\}.$$

Lemon and Lall then defined the related notion of weak team independence.

Definition 3. *We say that y is weakly team independent of x given s, and write $y \perp\!\!\!\perp x \mid s$ weakly, if*

$$p(x, y) \in \operatorname{conv}\{p(x, g \circ s) \mid g \in \mathcal{D}(S, Y)\}.$$

We can interpret the definitions of team independence and weak team independence in terms of the ability to generate simulated measurements (Lemon and Lall 2015). We have that y is team independent of x given s if we can use s to generate a simulated measurement $\hat{y}$ that has the same joint distribution with x and s as the original measurement y. In contrast, y is weakly team independent of x given s if we can use s to generate a simulated measurement $\hat{y}$ that has the same joint distribution with x as the original measurement y. (In particular, weak team independence does not require that $\hat{y}$ have the same joint distribution with s that y does.) The players are assumed to have access to common randomness when they generate these simulated measurements.

If y is team independent of x given s, then y must be weakly team independent of x given s. We can see this by marginalizing over s in the definition of team independence. More concretely, y is team independent of x given s if there exists $\theta \in \Delta(\mathcal{D}(S, Y))$ such that

$$p(x, y, s) = \sum_{g \in \mathcal{D}(S,Y)} \theta(g) p(x, g \circ s, s).$$

Summing both sides of this equation over s, we find that

$$\begin{aligned} p(x, y) &= \sum_{c \in S} p(x, y, s = c) \\ &= \sum_{c \in S} \sum_{g \in \mathcal{D}(S,Y)} \theta(g) p(x, g \circ s, s = c) \\ &= \sum_{g \in \mathcal{D}(S,Y)} \theta(g) p(x, g \circ s), \end{aligned}$$

and hence that y is weakly team independent of x given s. Lemon and Lall left it as an open problem to determine whether the converse result is true: that is, whether weak team independence implies team independence. After we define team sufficient statistics, we will present an example showing that weak team independence does not imply team independence.

Definition 4. *We say that s is a team statistic of y if there exists a function $h \in \mathcal{D}(Y, S)$ such that*

$$s(\omega) = h(y(\omega))$$

for all $\omega \in \Omega$.

Combining the definitions of team independence and team statistics gives the definition of team sufficient statistics.

Definition 5. *We say that s is a team sufficient statistic for x given y if s is a team statistic of y, and $y \perp\!\!\!\perp x \mid s$.*

Similarly, we define weak team sufficient statistics by combining the definitions of weak team independence and team statistics.

Definition 6. *We say that s is a weak team sufficient statistic for x given y if s is a team statistic of y, and $y \perp\!\!\!\perp x \mid s$ weakly.*

The following example shows that s being a weak team sufficient statistic does not guarantee that s is a team sufficient statistic. This also demonstrates that weak team independence does not imply team independence. Wu first used this example to demonstrate that team sufficient statistics do not enjoy the transitivity property of classical sufficient statistics (Wu 2013). (Recall that the transitivity property of classical sufficient statistics says that if $s^{(1)}$ is a statistic of $s^{(2)}$, and $s^{(1)}$ is a sufficient statistic, then $s^{(2)}$ must also be a sufficient statistic.)

Example 7. *Suppose z_1 and z_2 are independent Bernoulli random variables with success probability $1/2$. Consider a team decision problem with two agents, where $x = z_1$, $y_1 = (y_{11}, y_{12}) = (z_1, z_2)$, and $y_2 = (y_{21}, y_{22}) = (z_1, z_2)$. Consider the team statistic $s = (s_1, s_2)$ such that $s_1 = (s_{11}, s_{12}) = (y_{11}, y_{12})$ and $s_2 = y_{21}$. We claim that s is a weak team sufficient statistic, but not a team sufficient statistic for x given y. We can justify this result informally using the interpretation of sufficient statistics in terms of simulated measurements. We can use s and common randomness to generate a simulated measurement $\hat{y}$ that has the same joint distribution with x as the original measurement y as follows. We assume that our common randomness is a Bernoulli random variable z_3 with success probability $1/2$ that is independent of z_1 and z_2. Then we can take $\hat{y}_1 = (s_{11}, z_3)$*

and $\hat{y}_2 = (s_2, z_3)$. Thus, we can use s to generate a simulated measurement that has the same joint distribution with x, which implies that s is a weak team sufficient statistic.

However, we cannot use s to generate a simulated measurement $\hat{y}$ that has the same joint distribution with x and s as the original measurement. Then, in order for $\hat{y}_1$ to have the same joint distribution with s, we require that $\hat{y}_{12} = s_{12}$. Similarly, in order for $\hat{y}_2$ to have the same joint distribution with s, we require that $\hat{y}_{22} = \hat{y}_{12} = s_{12}$. However, agent 2 does not have access to the statistic s_{12}, so it is impossible to generate a simulated measurement that has the same joint distribution with s. This means that s is not a team sufficient statistic.

A more formal algebraic proof of our claim is as follows. We can represent the joint distribution of x and $y = (y_{11}, y_{12}, y_{21}, y_{22})$ using the vector $\phi \in \mathbf{R}^{32}$ such that

$$\begin{aligned} &p(x = a, y = b) \\ &\quad = \phi(16a + 8b_{11} + 4b_{12} + 2b_{21} + b_{22} + 1). \end{aligned}$$

Since (x, y) is uniformly distributed on the set

$$\{(0,0,0,0,0), (0,0,1,0,1), (1,1,0,1,0), (1,1,1,1,1)\},$$

we have that

$$\phi = \frac{1}{4}(e_1 + e_6 + e_{27} + e_{32}).$$

Consider the functions $g^{(1)}$ and $g^{(2)}$ defined in Tables 1 and 2. Using the same convention that we used to represent the joint distribution of x and y, we can represent the distributions of $(x, g^{(1)} \circ s)$ and $(x, g^{(2)} \circ s)$ by the vectors

$$\psi_1 = \frac{1}{2}(e_1 + e_{27}) \quad \textit{and} \quad \psi_2 = \frac{1}{2}(e_6 + e_{32}).$$

Then we have that

$$\phi = \frac{1}{2}\psi_1 + \frac{1}{2}\psi_2.$$

This proves that $p(x, y) \in \operatorname{conv}\{p(x, g \circ s) \mid g \in \mathcal{D}(S, Y)\}$, and hence that y is weakly team independent of x given s.

Now we show that y is not team independent of x given s. We have that $s = h(y)$, where $h(y) = ((y_{11}, y_{12}), y_{21})$. Lemon and Lall showed that s is a team sufficient statistic for x given y if and only if

$$p(x, y) \in \operatorname{conv}\left\{ p(x, g \circ s) \;\middle|\; \begin{array}{l} g \in \mathcal{D}(S, Y) \\ \textit{is a right inverse of } h \end{array} \right\}.$$

There are exactly four diagonal right inverses of h, which we denote $\gamma^{(k)}$ for $k = 1, 2, 3, 4$. We must have that $\gamma_1^{(k)}$ (the component of $\gamma^{(k)}$ corresponding to the first agent) is the identity map on $\{0, 1\}^2$; the components of $\gamma^{(k)}$ corresponding to the second agent are defined in Table 3. Using our usual convention for representing distributions on $X \times Y$, we can represent the joint distribution of x and $\gamma^{(k)} \circ s$ by the vector θ_k for $k = 1, 2, 3, 4$, where we define

$$\begin{aligned} \theta_1 &= \frac{1}{4}(e_1 + e_5 + e_{27} + e_{31}), \\ \theta_2 &= \frac{1}{4}(e_2 + e_6 + e_{27} + e_{31}), \\ \theta_3 &= \frac{1}{4}(e_1 + e_5 + e_{28} + e_{32}), \\ \theta_4 &= \frac{1}{4}(e_2 + e_6 + e_{28} + e_{32}). \end{aligned}$$

Observe that it is impossible to express ϕ as a convex combination of the θ_k. For example, we cannot put a positive weight on θ_1 because this results in positive components in the directions of e_5 and e_{31}, which are not found in ϕ. Similar arguments show that we cannot put positive weights on θ_2, θ_3, and θ_4. Thus, we see that ϕ cannot be written as a convex combination of the θ_k, which implies that s is not a team sufficient statistic for x given y.

Approximate Sufficient Statistics for Team Decision Problems

We have that y is team independent of x given s if $p(x, y)$ is contained in the convex hull of the set of distributions of the form $p(x, g \circ s)$, where $g \in \mathcal{D}(S, Y)$. It is natural to generalize this idea by considering the case when $p(x, y)$ is "close" to the convex hull of the set of distributions of the form $p(x, g \circ s)$. The following definition of approximate team independence makes this idea precise.

Definition 8. *We say that y is ϵ-approximately team independent of x given s with respect to the α-norm, and write $y \perp\!\!\!\perp_{\epsilon}^{\alpha} x \mid s$, if*

$$\min_{\theta \in \Delta(\mathcal{D}(S,Y))} \|p(x, y) - \tilde{p}_\theta(\tilde{x}, \tilde{g} \circ \tilde{s})\|_\alpha \le \epsilon,$$

where

$$\tilde{p}_\theta(\tilde{x}, \tilde{g} \circ \tilde{s}) = \sum_{g \in \mathcal{D}(S,Y)} \theta(g) p(x, g \circ s).$$

The following theorem bounds the suboptimality of using a policy based on s rather than y when y is ϵ-approximately team independent of x given s with respect to the α-norm. Note that s is not necessarily a statistic of y in this theorem: at this point in our analysis, y and s are simply two different sets of measurements available to the team, and we want to determine which set of measurements is better.

Theorem 9. *If $y \perp\!\!\!\perp_{\epsilon}^{\alpha} x \mid s$, then*

$$C^\star(s) \le C^\star(y) + \epsilon \|C(x, \phi \circ y)\|_{\alpha*}.$$

Proof. Suppose ϕ is an optimal policy of y, and θ achieves the minimum in the definition of an approximate sufficient statistic. Then we have that $\phi \circ g \in \mathcal{D}(S, U)$ for every $g \in \mathcal{D}(S, Y)$, and hence that

$$\begin{aligned} C^\star(s) &= \min_{\psi \in \mathcal{D}(S,U)} \mathbf{E}\, C(x, \psi \circ s) \\ &\le \mathbf{E}\, C(\tilde{x}, \phi \circ \tilde{g} \circ \tilde{s}). \end{aligned}$$

s_{11}	s_{12}	$g_1^{(1)}(s_1)$: $g_{11}^{(1)}(s_{11}, s_{12})$	$g_1^{(1)}(s_1)$: $g_{12}^{(1)}(s_{11}, s_{12})$	$g_1^{(2)}(s_1)$: $g_{11}^{(2)}(s_{11}, s_{12})$	$g_1^{(2)}(s_1)$: $g_{12}^{(2)}(s_{11}, s_{12})$
0	0	0	0	0	1
0	1	0	0	0	1
1	0	1	0	1	1
1	1	1	0	1	1

Table 1: definitions of $g^{(1)}$ and $g^{(2)}$ for the first agent

s_2	$g_2^{(1)}(s_2)$: $g_{21}^{(1)}(s_2)$	$g_2^{(1)}(s_2)$: $g_{22}^{(1)}(s_2)$	$g_2^{(2)}(s_2)$: $g_{21}^{(2)}(s_2)$	$g_2^{(2)}(s_2)$: $g_{22}^{(2)}(s_2)$
0	0	0	0	1
1	1	0	1	1

Table 2: definitions of $g^{(1)}$ and $g^{(2)}$ for the second agent

Using the definition of expectation, and the triangle inequality, we find that

$$\begin{aligned} &C^\star(s) - C^\star(y) \\ &\quad\le \mathbf{E}\, C(\tilde{x}, \phi \circ \tilde{g} \circ \tilde{s}) - \mathbf{E}\, C(x, \phi \circ y) \\ &\quad= \sum_{a\in X}\sum_{b\in Y} \tilde{p}_\theta(\tilde{x} = a, \tilde{g}\circ\tilde{s} = b) C(a, \phi^\star(b)) \\ &\qquad\quad - \sum_{a\in X}\sum_{b\in Y} p(x = a, y = b) C(a, \phi(b)) \\ &\quad\le \sum_{a\in X}\sum_{b\in Y} |\delta_\theta(a,b)||C(a, \phi(b))|, \end{aligned}$$

where we define

$$\delta_\theta(a,b) = \tilde{p}_\theta(\tilde{x} = a, \tilde{g}\circ\tilde{s} = b) - p(x = a, y = b).$$

Applying Hölder's inequality gives

$$C^\star(s) - C^\star(y) \le \|\delta_\theta\|_\alpha \|C(x, \phi\circ y)\|_{\alpha*}.$$

Because $y \perp\!\!\!\perp_\epsilon^\alpha x \mid s$, we have that

$$\|\delta_\theta\|_\alpha = \|p(x,y) - \tilde{p}_\theta(\tilde{x}, \tilde{g}\circ\tilde{s})\|_\alpha \le \epsilon,$$

and hence that

$$C^\star(s) \le C^\star(y) + \epsilon \|C(x, \phi\circ y)\|_{\alpha*}.$$

□

One of the most common metrics for evaluating the distance between probability distributions is the total-variation distance. The following corollary specializes Theorem 9 to a statement about the total-variation distance between $p(x,y)$ and $\tilde{p}_\theta(\tilde{x}, \tilde{g}\circ\tilde{s})$.

Corollary 10. *We have that*

$$C^\star(s) \le C^\star(y) + 2\,\mathrm{d_{TV}}(p(x,y), \mathcal{S})\|C\|_\infty,$$

where

$$\mathcal{S} = \mathrm{conv}\{p(x, g\circ s) \mid g \in \mathcal{D}(S, Y)\},$$

and

$$\mathrm{d_{TV}}(p(x,y), \mathcal{S}) = \min_{q\in\mathcal{S}} \mathrm{d_{TV}}(p(x,y), q(x,y)).$$

Proof. Suppose $q(x,y)$ is the distribution in $\mathcal{S}$ that is closest to $p(x,y)$. Applying Theorem 9 with $\alpha = 1$ gives

$$C^\star(s) \le C^\star(y) + \|p(x,y) - \tilde{p}(\tilde{x}, \tilde{g}\circ\tilde{s})\|_1 \|C(x, \phi\circ y)\|_\infty.$$

The total-variation distance and the 1-norm are related by the equation

$$\mathrm{d_{TV}}(p(x,y), \mathcal{S}) = \frac{1}{2}\|p(x,y) - \tilde{p}(\tilde{x}, \tilde{g}\circ\tilde{s})\|_1.$$

We also have that

$$\begin{aligned} \|C(x, \phi\circ y)\|_\infty &= \max_{a\in X, b\in y} |C(a, \phi(b))| \\ &\le \max_{a\in X, d\in U} |C(a, d)| \\ &= \|C\|_\infty. \end{aligned}$$

Combining these results, we find that

$$C^\star(s) \le C^\star(y) + 2\,\mathrm{d_{TV}}(p(x,y), \mathcal{S})\|C\|_\infty.$$

□

Computational Examples

Suppose we are given random variables x, y, and s, and a scalar $\alpha \in [1, \infty]$, and we want to evaluate the degree of approximate team independence of y and x given s with respect to the α-norm. In particular, we want to find the smallest ϵ such that y is ϵ-approximately team independent of x given s with respect to the α-norm. Let $\hat{p} \in \mathbf{R}^{X\times Y}$ be a vector representation of $p(x,y)$, and $\hat{M} \in \mathbf{R}^{(X\times Y)\times\mathcal{D}(S,Y)}$ be the matrix whose columns are corresponding vector representations of $p(x, g\circ s)$ for $g \in \mathcal{D}(S,Y)$. Then we want to solve the optimization problem

$$\begin{array}{ll} \underset{\theta\in\Delta(\mathcal{D}(S,Y))}{\text{minimize:}} & \|\hat{M}\theta - \hat{p}\|_\alpha \\ \text{subject to:} & \sum_{g\in\mathcal{D}(S,Y)} \theta(g) = 1, \\ & \theta(g) \ge 0, \quad g \in \mathcal{D}(S,Y). \end{array}$$

This is a convex optimization problem that can be solved efficiently.

A single-player example

First, we give some results on approximate sufficient statistics for a family of classical single-player decision problem. Suppose the random variable x is uniformly distributed on the set $\{1, \dots, n\}$, and, given x, the measurements $y_1, \dots, y_m$ are conditionally independent and identically distributed uniformly on the set $\{1, \dots, x\}$. It is a well-known classical result that a sufficient statistic for x given

	$\gamma_2^{(1)}(s_2)$		$\gamma_2^{(2)}(s_2)$		$\gamma_2^{(3)}(s_2)$		$\gamma_2^{(4)}(s_2)$	
s_2	$\gamma_{21}^{(1)}(s_2)$	$\gamma_{22}^{(1)}(s_2)$	$\gamma_{21}^{(2)}(s_2)$	$\gamma_{22}^{(2)}(s_2)$	$\gamma_{21}^{(3)}(s_2)$	$\gamma_{22}^{(3)}(s_2)$	$\gamma_{21}^{(4)}(s_2)$	$\gamma_{22}^{(4)}(s_2)$
0	0	0	0	1	0	0	0	1
1	1	0	1	0	1	1	1	1

Table 3: definitions of the $\gamma^{(k)}$ for the second player

	$n=3$		$n=4$	
k	$m=1$	$m=2$	$m=1$	$m=2$
2	0.0370	0.1562	0.0457	0.2230
3			0.0139	

Table 4: approximation levels for conditionally uniform observations

	$n=3$		$n=4$	
k	$m=1$	$m=2$	$m=1$	$m=2$
2	$\{1\}$	$\{(1,1)\}$	$\{1\}$	$\{(1,1)\}$
3			$\{1\},\{2\}$	

Table 5: approximate sufficient statistics for conditionally uniform observations

$y=(y_1,\ldots,y_m)$ is $s_0=\max(y_1,\ldots,y_m)$. Thus, we can make optimal decisions by storing which of the n possible values that we observe for s_0. However, suppose we want to compress our measurements even further, so that our statistic s takes on at most k distinct values. The approximation levels (measured using the total-variation distance) for various values of the problem parameters are given in Table 4. We can think of a statistic as a partition of the set of measurements. Moreover, we can represent a partition by listing all but one of the sets that form the partition (the last set being omitted for convenience). The optimal approximate sufficient statistics are described in this way in Table 5.

A multi-agent example

Now we present numerical results about the approximate sufficiency levels of different statistics for a team decision problem with respect to different norms. We assume that the state x is uniformly distributed on $\{1,2\}$. Let z_1, z_2, and z_3 be conditionally independent given x, and let the conditional distribution of z_j given x be uniformly distributed on $\{1,x\}$. The measurements of agent j are $y_j=(z_j,z_3)$ for $j=1,2$. We consider three different sets of statistics.

1. In the first set of statistics, each agent stores the common measurement: that is, $s_1^{(1)}=y_{12}$ and $s_2^{(1)}=y_{22}$.
2. In the second set of statistics, each agent stores its private measurement: that is, $s_1^{(2)}=y_{11}$ and $s_2^{(2)}=y_{21}$.
3. In the third set of statistics, the first agent stores its private measurement, and the second agent stores the common measurement: that is, $s_1^{(3)}=y_{11}$ and $s_2^{(3)}=y_{22}$.

	statistics		
norm	$s^{(1)}$	$s^{(2)}$	$s^{(3)}$
1-norm	0.3750	0.3750	0.3750
2-norm	0.1793	0.1168	0.1168
∞-norm	0.1250	0.0500	0.0500
weighted 1-norm	0.5625	0.5000	0.5000

Table 6: approximate sufficiency levels for different statistics with respect to different norms

Intuitively, the team is given three measurements of the state x: one private measurement for each agent, and one common measurement shared by both agents. Each player can only store one measurement, and must decide whether to store the common measurement or the private measurement. We consider the cases when both players store the common measurement, both players store their private measurements, and one player stores its private measurement while the other player stores the common measurement.

We give the approximate sufficiency levels for these statistics with respect to various norms in Table 6. In addition to the usual 1-, 2-, and ∞-norms, we include a weighted 1-norm given by

$$\sum_{a\in X}\sum_{b\in Y}(1+[y_{12}=y_{22}])|p(x,y)|,$$

where

$$[y_{12}=y_{22}]=\begin{cases}1 & y_{12}=y_{22},\\ 0 & \text{otherwise}\end{cases}$$

is the Iverson bracket for the statement $y_{12}=y_{22}$. This weighted norm puts extra emphasis on ensuring that the simulated measurements satisfy $\hat{y}_{12}=\hat{y}_{22}$. One might think that the agents should both store the common measurement because coordination among agents is often important in team decision problems. However, note that coordination is also achieved using the common randomness that forms the basis of a randomized team policy. Therefore, the team is actually able to do better by storing as many observations as possible. This explains why the agents should store different measurements (either both agents should store their private measurements, or one should store its private measurement, and the other should store the common measurement).

Summary and Future Work

Weak team sufficient statistics are necessary and sufficient for making optimal decisions, and may allow us to solve complex team decision problems by significantly reducing the size of the search space. The generalization of these statistics to approximate sufficient statistics may allow us to sacrifice exact optimality for even greater reductions in the size of the search space. There are still many open questions regarding weak and approximate team sufficient statistics. In particular, we need a calculus for these statistics that is similar to the theory used to construct team sufficient statistics from single-player statistics (Wu 2013). Additional work is needed to extend weak and approximate team sufficient statistics to dynamic problems, and to determine the relationship between these statistics and the notion of probabilistic equivalence used in the Dec-POMDP community (Dibangoye et al. 2013; Oliehoek, Whiteson, and Spaan 2009; Oliehoek 2013).

Acknowledgments

This research was partially supported by AFOSR MURI grant FA9550-10-1-0573.

References

Bernstein, D. S.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research* 27(4):819 – 840.

Dibangoye, J. S.; Amato, C.; Buffet, O.; and Charpillet, F. 2013. Optimally solving dec-POMDPs as continuous-state MDPs. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, 302 – 308. AAAI Press.

Hansen, E.; Bernstein, D.; and Zilberstein, S. 2004. Dynamic programming for partially observable stochastic games. In *AAAI Spring Symposium*, volume 4, 709 – 715. AAAI.

Ho, Y.-C., and Chu, K. C. 1972. Team decision theory and information structures in optimal control problems – part I. *IEEE Transactions on Automatic Control* 17(1):15 – 22.

Kube, C. R., and Zhang, H. 1997. Task modelling in collective robotics. In *Autonomous Robots*, volume 4. Springer. 53 – 72.

Lemon, A., and Lall, S. 2015. Weak sufficient statistics for team decision problems. In *54th IEEE Conference on Decision and Control*, To Appear. IEEE.

Lessard, L., and Lall, S. 2015. Optimal control of two-player systems with output feedback. *IEEE Transactions on Automatic Control* 60(8):2129 – 2144.

Marschak, J. 1955. Elements for a theory of teams. *Management Science* 1(2):127–137.

Oliehoek, F. A.; Whiteson, S.; and Spaan, M. T. J. 2009. Lossless clustering of histories in decentralized POMDPs. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems*, 577 – 584. International Foundation for Autonomous Agents and Multiagent Systems.

Oliehoek, F. A. 2013. Sufficient plan-time statistics for decentralized POMDPs. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, 302 – 308. AAAI Press.

Papadimitriou, C. H., and Tsitsiklis, J. N. 1987. The complexity of markov decision processes. *Mathematics of Operations Research* 12(3):441 – 450.

Radner, R. 1962. Team decision problems. *The Annals of Mathematical Statistics* 33(3):857 – 881.

Tsitsiklis, J. N., and Athans, M. 1985. On the complexity of decentralized decision making and detection problems. *IEEE Transactions on Automatic Control* 30(5):440 – 446.

Wu, J., and Lall, S. 2010a. A dynamic programming algorithm for decentralized Markov decision processes with a broadcast structure. In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, 6143 – 6148. IEEE.

Wu, J., and Lall, S. 2010b. An exact dynamic programming solution for a decentralized two-player markov decision process. In *Proceedings of the AAAI Spring Symposium*, 112 – 119. AAAI.

Wu, J., and Lall, S. 2014a. Sufficient statistics for multi-agent decision problems. In *52nd Annual Allerton Conference on Communication, Control, and Computing*, 467 – 474. IEEE.

Wu, J., and Lall, S. 2014b. A theory of sufficient statistics for teams. In *53rd Annual Conference on Decision and Control*, 2628 – 2635. IEEE.

Wu, J., and Lall, S. 2015. Sufficient statistics for dynamic teams. In *American Control Conference (ACC)*, 5465 – 5470. IEEE.

Wu, J. 2013. *Sufficient Statistics for Team Decision Problems*. Ph.D. Dissertation, Stanford University.