You can’t save all the pandas: impossibility results for privacy-preserving tracking

Yulin Zhang and Dylan A. Shell

Department of Computer Science and Engineering,
Texas A&M University,
College Station, TX 77840, USA
yulinzhang|dshell@tamu.edu

Abstract. We consider the problem of target tracking whilst simultaneously preserving the target’s privacy as epitomized by the panda tracking scenario introduced by O’Kane at *WAFR’08*. The present paper reconsiders his formulation, with its elegant illustration of the utility of ignorance, and the tracking strategy he proposed, along with its completeness. We explore how the capabilities of the robot and panda affect the feasibility of tracking with a privacy stipulation, uncovering intrinsic limits, no matter the strategy employed. This paper begins with a one-dimensional setting and, putting the trivially infeasible problems aside, analyzes the strategy space as a function of problem parameters. We show that it is not possible to actively track the target as well as protect its privacy for every nontrivial pair of tracking and privacy stipulations. Secondly, feasibility is sensitive, in several cases, to the information available to the robot at the start — conditions we call initial I-state dependent cases. Quite naturally in the one-dimensional model, one may quantify sensing power by the number of perceptual (or output) classes available to the robot. The number of initial I-state dependent conditions does not decrease as the robot gains more sensing power and, further, the robot’s power to achieve privacy-preserving tracking is bounded, converging asymptotically with increasing sensing power. Finally, to relate some of the impossibility results in one dimension to their higher-dimensional counterparts, including the planar panda tracking problem studied by O’Kane, we establish a connection between tracking dimensionality and the sensing power of a one-dimensional robot.

1 Introduction

Most roboticists see uncertainty as something which should be minimized or even eliminated if possible. But, as robots become widespread, it is likely that there will be a shift in thinking— robots that know too much are also problematic in their own way. A robot operating in your home, benignly monitoring your activities and your daily routine, for example to schedule vacuuming at unobtrusive times, possesses information that is valuable. There are certainly those who could derive profit from it. A tension exists between information necessary for a robot to be useful and information which could be sensitive if mistakenly disclosed or stolen. But the problem of establishing the minimal information required to perform a particular task and the problem of analyzing the trade-off between information and performance, despite both being fundamental challenges, remain largely uncharted territory — notwithstanding [1] and [2].

The present paper's focus is on the problem of tracking a target whilst preserving the target's privacy. Though fairly narrow, this is a crisply formulated instance of the broader dilemma of balancing the information a robot possesses: the robot must maintain some estimate of the target's pose, but information that is too precise is an unwanted intrusion and potential hazard if leaked. The setting we examine, the *panda tracker problem*, is due to O'Kane [3] who expressed the idea of uncertainty being valuable and aloofness deserving respect. The following, quoted verbatim from [3, p. 1], describes the scenario:

> "A giant panda moves unpredictably through a wilderness preserve. A mobile robot tracks the panda's movements, periodically sensing partial information about the panda's whereabouts and transmitting its findings to a central base station. At the same time, poachers attempt to exploit the presence of the tracking robot—either by eavesdropping on its communications or by directly compromising the robot itself—to locate the panda. We assume, in the worst case, that the poachers have access to any information collected by the tracking robot, but they cannot control its motions. The problem is to design the tracking robot so that the base station can record coarse-grained information about the panda's movements, without allowing the poachers to obtain the fine-grained position information they need to harm the panda."

Note that it is not sufficient for the robot to simply forget or to degrade sensor data via post-processing because the adversary may have compromised these operations, simply writing the information to separate storage.

Before formalizing our approach to the problem, which we do in detail in Section 2, we give an overview of the questions of interest to us. We also give an outline of the structure of the paper and summarize the reported results.

One can view the informational constraints as bounds: (1) A maximal- or upper-bound specifies how coarse the tracking information can be. The robot is not helpful in assuring the panda's well-being when this bound is exceeded. (2) A second constraint, a lower-bound, stipulates that if the information is more fine-grained than some threshold, a poacher may succeed in having his wicked way. The problem is clearly infeasible when the lower-bound exceeds the upper-bound. What of other circumstances? Is it always possible to ensure that one will satisfy both bounds indefinitely? In the original paper, O'Kane [3] proposed a tracking strategy for a robot equipped with a two-bit quadrant sensor, which certainly succeeds in several circumstances. As no claim of completeness was made, will the strategy work for all feasible bounds? And how are the strategies affected by improvements in the sensing capabilities of the robot?

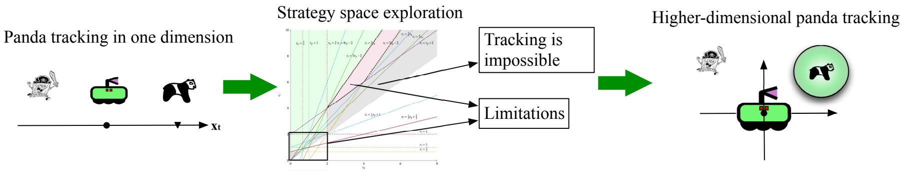
Panda tracking in one dimension
x_t
Strategy space exploration
Tracking is impossible
Limitations
Higher-dimensional panda tracking

Fig. 1: An overview of the approach taken in the paper.

Our examination of these questions follows the outline shown in Figure 1. We start with a one-dimensional panda tracking problem, analyzing it in detail in Section 3, and find that it is impossible for the robot to achieve privacy-preserving tracking for all nontrivial tracking and privacy stipulations. In addition, there are instances where privacy-preserving tracking is limited, depending crucially on the initial information available to the robot. Then, in Section 4, we show that the impossibility conclusion holds for O'Kane's original planar panda tracking problem too.

2 Problem description: one-dimensional panda tracking

The original problem was posed in two dimensions with the robot and panda inhabiting a plane that is assumed to be free of obstacles. They move in discrete time-steps, interleaving their motions. A powerful adversary, who is interested in computing the possible locations of the panda, is assumed to have access to the full history of information. Any information is presumed to be used optimally by the adversary in reconstruction of possible locations of the panda — by which we mean that the region that results is both sound (that is, consistently accounted for by the information) but is also tight (no larger than necessary). The problem is formulated without needing to appeal to probabilities by considering only worst-case reasoning and by using motion- and sensor-models characterized by regions and applying geometric operations.

Information stipulation: The tracking and privacy requirements were specified as two disks. The robot is constrained to ensure that the region describing possible locations of the panda always fits inside the *tracking disk*, which has the larger diameter of the two. The *privacy disk* imposes the requirement that it always be possible to place the smaller disk wholly inside the set of possible locations.

Sensor model: As reflected in the title of his paper, O'Kane considered an unconventional sensor that outputs only two bits of information per measurement. With origin centered on the robot, the sensor outputs the quadrant containing the panda.

Target motion model: The panda moves unpredictably with bounded velocity. Newly feasible locations can be modeled as the convolution of the previous time-step's region with a disk, sized appropriately for the time-step interval.

Now, by way of simplification, consider pandas and robots that inhabit obstacle-free one-dimensional worlds, each only moving left or right along a line.

Information stipulation: Using the obvious 1-dimensional analogue, now the robot tracker has to bound its belief about the panda's potential locations to an interval of minimum size r_p (p for privacy) and maximum size r_t (t for tracking).

Sensor model: Most simply, the quadrant sensor corresponds to a one-bit sensor indicating whether the panda is on the robot's left- or right-hand side. When the robot is at u_1, the sensor indicates whether the panda is within $(-\infty, u_1]$ or (u_1, ∞).

We have sought a way to explore how modifying the robot's sensing capabilities alters its possible behavior. Our approach is to give the robot m set-points $u_1 < u_2 < \cdots < u_m$, each within the robot's control, so that it can determine which of the $m+1$ non-overlapping intervals contains the panda. With m set-points one can model any sensor that fully divides the state space and produces at most $m+1$ observations.* The case with $m=1$ is the straightforward analogue of the quadrant sensor. Increasing m yields a robot with greater sensing power, since the robot has a wider choice of how observations should inform itself of the panda's location.

Target motion model: The convolution becomes a trivial operation on intervals.

Figure 2 is a visual example with $m=2$. The panda is sensed by the robot as falling within one of the following intervals: $(-\infty, u_1]$, $(u_1, u_2]$, (u_2, ∞), where $u_1, u_2 \in \mathbb{R}$ and $u_1 < u_2$. And these three intervals are represented by observation values: 0, 1 and 2. For simplicity, no constraints are imposed on the robot's motion and we assume that each time-step, the robot can pick positions of $u_1 < \cdots < u_m$ as it likes.

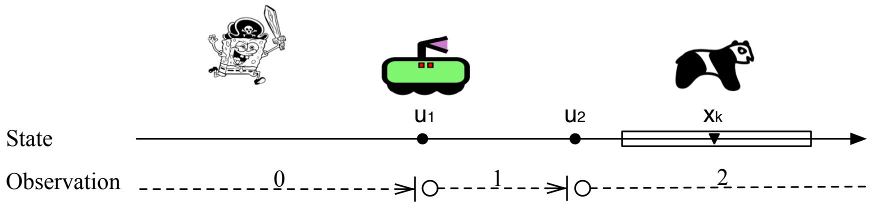

Fig. 2: Panda tracking in one dimension. (Inimitable artwork for the robot and panda adapted from the original in [3, p. 2].)

Notation and model:

The 1-dim. problem is formulated as follows. The panda's location is represented as a single coordinate indexed by discrete time. At stage k the location of the panda is $x_k \in \mathbb{R}$. Incorporating accumulated knowledge about the panda's possible location, sensor readings, and the movement model permits the robot (and adversary) to maintain knowledge of the panda's possible location after it moves and between sensing steps. The set of conceivable locations of the panda (a subset of $\mathbb{R}$) is a geometric realization of an information-state or I-state, as formalized and treated in detail by LaValle [4]. In this paper, we take the I-state as an interval.

The movement of the panda per time-step is bounded by length $\frac{\delta}{2}$, meaning that the panda can move at most $\frac{\delta}{2}$ in either direction. We use η_k to denote the robot's knowledge of the panda after the observation taken at time k. In evolving from η_k to η_{k+1} the robot's I-state first transits to an intermediate I-state, which we write η^-_{k+1} representing the state after adding the uncertainty arising from the panda's movement, but before observation $k+1$. Since this update occurs before the sensing information is incorporated, we refer to η^-_{k+1} as the *prior* I-state for time $k+1$. Updating I-state η_k involves mapping every $x_k \in \eta_k$ to $[x_k - \frac{\delta}{2}, x_k + \frac{\delta}{2}]$, the resultant I-state, η^-_{k+1}, being the union of the results.

*This is not a model of any physical sensor of which we are aware. The reader, finding this too contrived, may find merit in the choice later, e.g., for n-dimensional tracking (see Lemma 6).

Sensor reading updates to the I-state depend on the values of $u_1(k), u_2(k), \ldots, u_m(k)$, which are under the control of the robot. The sensor reports the panda's location to within one of the $m+1$ non-empty intervals: $(-\infty, u_1(k)], (u_1(k), u_2(k)], (u_2(k), u_3(k)], \ldots, (u_m(k), \infty)$. If we represent the observation at time k as a non-empty interval $y(k)$ then the *posterior* I-state η_k is updated as $\eta_k = \eta_k^- \cap y(k)$.

For every stage k, the robot chooses a sensing vector $\mathbf{v_k} = [u_1(k), u_2(k), \ldots, u_m(k)]$, $u_i(k) < u_j(k)$ if $i < j$, $u_i(k) \in \mathbb{R}$, so as to achieve the following conditions:

1. *Privacy Preserving Condition (PPC)*: The size of any I-state $\eta_k = [a,b]$ should be at least r_p. That is, for every stage k, $|\eta_k| = b - a \geq r_p$.
2. *Target Tracking Condition (TTC)*: The size of any I-state $\eta_k = [a,b]$ should be at most r_t. That is, for every stage k, $|\eta_k| = b - a \leq r_t$.

3 Privacy-preserving tracking

Given specific problem parameters, we are interested in whether there is always some $\mathbf{v_k}$ that a robot can select to track the panda while satisfying the preceding conditions.

Definition 1. *A 1-dim. panda tracking problem is a tuple $P_1 = (\eta_0, r_p, r_t, \delta, m)$, in which*
1) the initial I-state η_0 describes all the possible initial locations of the panda;
2) the privacy bound r_p gives a lower bound on the I-state size;
3) the tracking bound r_t gives a upper bound on the I-state size;
4) parameter δ describes the panda's (fastest) motion;
5) the sensor capabilities are given by the number m.

Definition 2. *The 1-dim. panda tracking problem $P_1 = (\eta_0, r_p, r_t, \delta, m)$ is privacy preservable, written as predicate $\boldsymbol{PP}(P_1)$, if starting with $|\eta_0| \in [r_p, r_t]$, there exists some strategy π to determine a $\mathbf{v_k}$ at each time-step, such that the Privacy Preserving Condition holds forever. Otherwise, the problem P_1 is not privacy preservable: $\neg\boldsymbol{PP}(P_1)$.*

Definition 3. *The 1-dim. panda tracking problem $P_1 = (\eta_0, r_p, r_t, \delta, m)$ is target trackable, $\boldsymbol{TT}(P_1)$, if starting with $|\eta_0| \in [r_p, r_t]$, there exists some strategy π to determine a $\mathbf{v_k}$ at each time-step, such that the Target Tracking Condition holds forever. Otherwise, the problem P_1 is not target trackable: $\neg\boldsymbol{TT}(P_1)$.*

To save space, we say a problem P_1 and also its strategy π are **PP** if $\mathbf{PP}(P_1)$. Similarly, both P_1 and its strategy π will be called **TT** if $\mathbf{TT}(P_1)$. Putting aside trivially infeasible P_1 where $r_p > r_t$, we wish to know which problems are both **PP** and **TT**. Next, we explore the parameter space to classify the various classes of problem instances by investigating the existence of strategies.

3.1 Roadmap of technical results

The results follow from several lemmas. The roadmap in Figure 3 provides a sense of how the pieces fit together to help the reader keep an eye on the broader picture.

Our approach begins, first, by dividing the strategy space into 'teeth' and 'gaps' according to the speed (or size) of the panda's motion. **PP** and **TT** tracking strategies

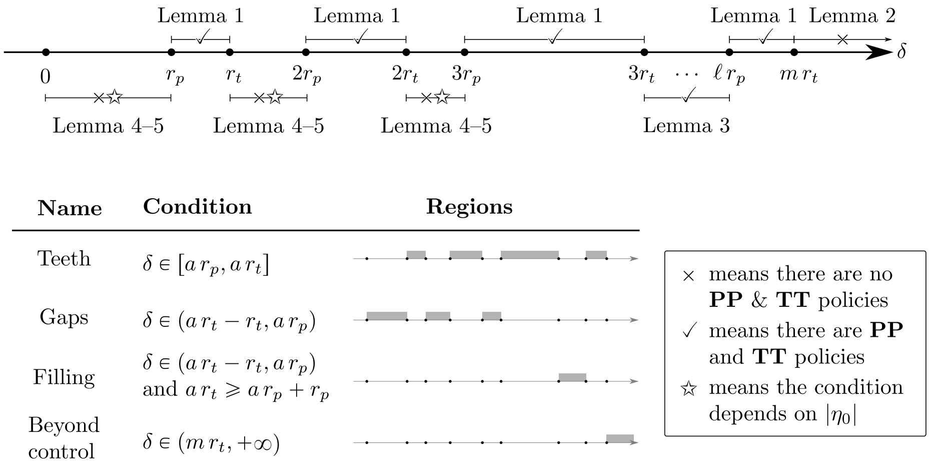

Lemma 1
Lemma 1
Lemma 1
Lemma 1
Lemma 2
δ
0
r_p
r_t
$2r_p$
$2r_t$
$3r_p$
$3r_t$
$\cdots$
ℓr_p
$m r_t$
Lemma 4–5
Lemma 4–5
Lemma 4–5
Lemma 3
Name
Condition
Regions
Teeth
$\delta \in [a\, r_p, a\, r_t]$
Gaps
$\delta \in (a\, r_t - r_t, a\, r_p)$
Filling
$\delta \in (a\, r_t - r_t, a\, r_p)$
and $a\, r_t \geqslant a\, r_p + r_p$
Beyond control
$\delta \in (m\, r_t, +\infty)$
× means there are no **PP** & **TT** policies
✓ means there are **PP** and **TT** policies
☆ means the condition depends on $|\eta_0|$

Fig. 3: Roadmap of results for 1-dim. panda tracking with *m* sensing parameters.

are shown to exist in teeth regions and the last gap region for all initial I-states that satisfy $|\eta_0| \in [r_p, r_t]$. The remaining regions differ: their feasibility depends not only on the panda's motion but also on the size of initial I-state. Dealing with the remaining gap regions is simplified considerably by focusing on $\mathbf{v_k}$ that divides η_k^- evenly. We use $s(i)$ to represent the choice of evenly dividing the prior I-state interval into i parts (the mnemonic is *s* for *split*). In practice, we will usually compare choices $s(a)$ and $s(a+1)$, for positive integer constant a determined from the speed of the panda's motion ($\frac{\delta}{2}$).

The action of evenly dividing the I-state interval η_k^- via $s(\cdot)$ is useful because, after the sensor reading has been processed, one knows the degree of the uncertainty involved, i.e., the size of the I-state interval η_k. An $s(i)$ action results in greater post-sensing uncertainty than an $s(i+1)$ does. Strategies consisting only of $s(i)$ and $s(i+1)$ actions enable examination of the evolution of interval sizes. If the I-state resulting from an $s(i)$ violates the tracking constraint, other actions dividing the I-state into at most i parts, though they be uneven parts, will also violate the tracking bound because we must guarantee a solution no matter which interval the panda happens to be in. Analogously, the $s(i+1)$ case is often useful in analyzing violation of the privacy bound.

We have found it helpful to think of the problem as playing a multi-stage game against suicidal pandas who strategically choose the movement that tends to violate the bounds. The robot's job is to save all of them. According to whether an $s(i)$ or $s(i+1)$ may be performed in the I-state or not, we are able to divide the remaining strategy space into four cases where two of them are non-**PP** or non-**TT**, one is **PP** and **TT** regardless of the initial I-state, and the other is I-state dependent.

3.2 Main results

In this section, we follow the roadmap in Figure 3.

Lemma 1. *For any 1-dim. panda tracking problem $P_1 = (\eta_0, r_p, r_t, \delta, m)$, if $\delta \in [ar_p, ar_t]$, where $a \in \mathbb{Z}^+$, $a \leq m$, then $\boldsymbol{PP}(P_1) \wedge \boldsymbol{TT}(P_1)$.*

Proof. A **PP** and **TT** strategy is given in this proof. For any $|\eta_k| \in [r_p, r_t]$ the prior I-state has size $|\eta_{k+1}^-| = |\eta_k| + \delta$. Since $\delta \in [ar_p, ar_t](a \leq m)$, $|\eta_{k+1}^-| \in [ar_p + r_p, ar_t + r_t]$. By taking action $s(a+1)$, which is possible since $a \leq m$, we get $|\eta_{k+1}| = \frac{1}{a+1}|\eta_{k+1}^-| \in [r_p, r_t]$. That is, if $|\eta_0| \in [r_p, r_t]$ and we take action $s(a+1)$, then $|\eta_k| \in [r_p, r_t]$. Therefore, there exists a strategy (always take action $s(a+1)$) for P_1, so that the privacy-preserving tracking conditions *PPC* and *TTC* are always both satisfied when $\eta_0 \in [r_p, r_t]$. □

Lemma 2. *For any 1-dim. panda tracking problem* $P_1 = (\eta_0, r_p, r_t, \delta, m)$*, if* $\delta \in (mr_t, \infty)$*, then* $\neg \boldsymbol{TT}(P_1)$.

Proof. The tracking stipulation is proved to be violated eventually. Given the constraint of m sensing parameters, the prior I-state $|\eta_{k+1}^-| = |\eta_k| + \delta$ can be divided into at most $m+1$ parts. Among these $m+1$ posterior I-states, if m of them reach the maximum size r_t, the size of the remaining I-state is $|\eta_k| + \delta - mr_t$. If none of the resulting I-states violate the tracking bound, then the size of the smallest resulting I-state is $|\eta_k| + \delta - mr_t$. Since $\delta > mr_t$, the size of the smallest resulting I-state must increase by some positive constant $\delta - mr_t$. After $\lceil \frac{r_t - r_p}{\delta - mr_t} \rceil$ stages, the I-state will exceed r_t. So it is impossible to ensure that the tracking bound will not eventually be violated. □

Lemma 3. *For 1-dim. panda tracking problem* $P_1 = (\eta_0, r_p, r_t, \delta, m)$*, if* $\delta \in (ar_t - r_t, ar_p)$*, where* $a \in \mathbb{Z}^+$*,* $a \leq m$ *and* $ar_t \geq ar_p + r_p$*, then* $\boldsymbol{PP}(P_1) \wedge \boldsymbol{TT}(P_1)$.

Proof. A **PP** and **TT** strategy is given in this proof. Since $\delta \in (ar_t - r_t, ar_p)$ and $ar_t > ar_p + r_p$, $|\eta_{k+1}^-| = |\eta_k| + \delta \in (ar_p, ar_t + r_t) \subset L_1 \cup L_2$, where $L_1 = [ar_p, ar_t]$ and $L_2 = [ar_p + r_p, ar_t + r_t]$. If action $s(a)$ is performed when $|\eta_{k+1}^-| \in L_1$ and $s(a+1)$ is performed when $|\eta_{k+1}^-| \in L_2$, the resulting I-state satisfies $|\eta_{k+1}| \in [r_p, r_t]$. Hence, there is a strategy consisting of $s(a)$ and $s(a+1)$, for the problem P_1 such that the *PPC* and *TTC* are always both satisfied when $\eta_0 \in [r_p, r_t]$. □

Lemma 4. *For any 1-dim. panda tracking problem* $P_1 = (\eta_0, r_p, r_t, \delta, m)$*, if* $\delta \in (ar_t - r_t, ar_p)$*, where* $a \in \mathbb{Z}^+$*,* $a \leq m$*,* $ar_t < ar_p + r_p$*, then* $\neg \boldsymbol{PP}(P_1) \vee \neg \boldsymbol{TT}(P_1)$ *when either: (i)* $r_p > ar_t - \delta$ *or (ii)* $r_t < (a+1)r_p - \delta$.

Proof. In case (i), $r_p > ar_t - \delta$, so the size of the prior I-state satisfies $|\eta_{k+1}^-| = |\eta_k| + \delta > r_p + \delta > ar_t$. That is, if η_{k+1}^- is divided into at most a parts, the largest posterior I-state η_{k+1} will violate the tracking stipulation at the next time-step. Thus, the prior I-state η_{k+1}^- must be divided into at least $a+1$ parts. But by dividing $|\eta_{k+1}^-|$ into at least $a+1$ parts, among all the resulting posterior I-states, it can be shown that the smallest I-state will eventually violate the privacy stipulation. The smallest size of the resulting posterior I-state $|\eta_{k+1}|_{smallest}$ is no greater than the average size $\frac{|\eta_k| + \delta}{a+1}$, when dividing η_{k+1}^- into $a+1$ parts. For the smallest posterior I-state, the decrease in size is $\Delta_- = |\eta_k| - |\eta_{k+1}|_{smallest} \geq |\eta_k| - \frac{|\eta_k| + \delta}{a+1} \geq \frac{ar_p - \delta}{a+1} > 0$. Hence, eventually after $\lceil \frac{(a+1)(r_t - r_p)}{ar_p - \delta} \rceil$ steps, the smallest I-state will violate the privacy stipulation and put the panda in danger. Similarly, if the prior I-state is divided into less than $a+1$ parts, the average size will become larger and the largest posterior I-state will violate the tracking stipulation, as it must increase at least as much as the average.

The same conclusion is reached for (ii), when $r_t < (a+1)r_p - \delta$, along similar lines. □

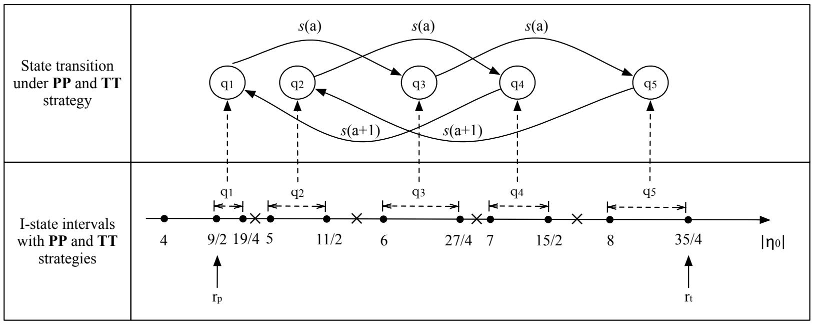
State transition under PP and TT strategy
s(a)
s(a)
s(a)
q1
q2
q3
q4
q5
s(a+1)
s(a+1)
I-state intervals with PP and TT strategies
q1
q2
q3
q4
q5
4
9/2
19/4
5
11/2
6
27/4
7
15/2
8
35/4
|η0|
rp
rt

Fig. 4: A problem instance with $r_p = \frac{9}{2}, r_t = \frac{35}{4}, \delta = 2, a = 1$, where there exists a **PP** and **TT** tracking strategy for $|\eta_0| \in q_1 \cup q_2 \cup q_3 \cup q_4 \cup q_5$.

Lemma 5. *The tracking and privacy-preserving properties of a 1-dim. panda tracking problem* $P_1 = (\eta_0, r_p, r_t, \delta, m)$ *depend on the size of the initial I-state* $|\eta_0|$*, if* $\delta \in (ar_t - r_t, ar_p)$*, where* $a \in \mathbb{Z}^+$*,* $a \leq m$ *and* $r_p \leq ar_t - \delta < (a+1)r_p - \delta \leq r_t$*.*

Proof. The proof proceeds by showing, firstly, that there are certain initial I-states for which the problem is not both **PP** and **TT**. Next, we show that there exist other initial I-states where there are strategies which satisfy the conditions for the problem to be **PP** and **TT**.

First, for $|\eta_0| \in (ar_t - \delta, (a+1)r_p - \delta)$, if the prior I-state is divided into at most a parts, the size of the largest posterior I-state at the next time-step ($t = 1$) is $|\eta_1| = \frac{|\eta_1^-|+\delta}{a} > \frac{ar_t-\delta+\delta}{a} = r_t$, which will violate the tracking stipulation. If the prior I-state is divided into at least $a+1$ parts, the size of posterior I-state at the next step ($t = 1$) is $|\eta_1| = \frac{|\eta_1^-|+\delta}{a+1} < \frac{(a+1)r_p-\delta}{a+1} = r_p$, which will violate the privacy stipulation. Hence, there are no strategies, which are both **PP** and **TT** if $|\eta_0| \in (ar_t - \delta, (a+1)r_p - \delta)$.

Second, we provide an example strategy that serves as an existence proof for **PP** and **TT** instances with $|\eta_0| \in [r_p, ar_t - \delta) \cup ((a+1)r_p - \delta, r_t]$. Consider the instance $r_p = \frac{9}{2}, r_t = \frac{35}{4}, \delta = 2, a = 1$, where there are **PP** and **TT** strategies for $\eta_0 \in q_1 \cup q_2 \cup q_3 \cup q_4 \cup q_5$, where $q_1 = [\frac{9}{2}, \frac{19}{4}], q_2 = [5, \frac{11}{2}], q_3 = [6, \frac{27}{4}], q_4 = [7, \frac{15}{2}], q_5 = [8, \frac{35}{4}]$. By always taking $s(a)$ and $s(a+1)$, the resulting I-state remains within $q_1, q_2, \ldots, q_5$ as shown in Figure 4. That is, there are strategies for those I-states in $q_1 \cup q_2 \cup \cdots \cup q_5$ and the problem is **PP** and **TT** under such circumstances. □

Next, we collect the low-level results presented thus far to clarify their interplay.

3.3 Aggregation and summary of results

To have a clearer sense of how these pieces fit together, we found it helpful to plot the space of problem parameters and examine how the preceding theorems relate visually. Figure 5 contains subfigures for increasingly powerful robots (in terms of sensing) with $m = 1, 2, 3, 4$. The white regions represent the trivial $r_p > r_t$ instances; otherwise the whole strategy space is categorized into the following subregions: **PP** and **TT** strategy space (colored green), not both **PP** and **TT** space (colored gray), and regions dependent on the initial I-states (colored pink — some caution is advised as particular values of η_0 are not visible in this projection). When summarized in this way, the results permit examination of how sensor power affects the existence of solutions.

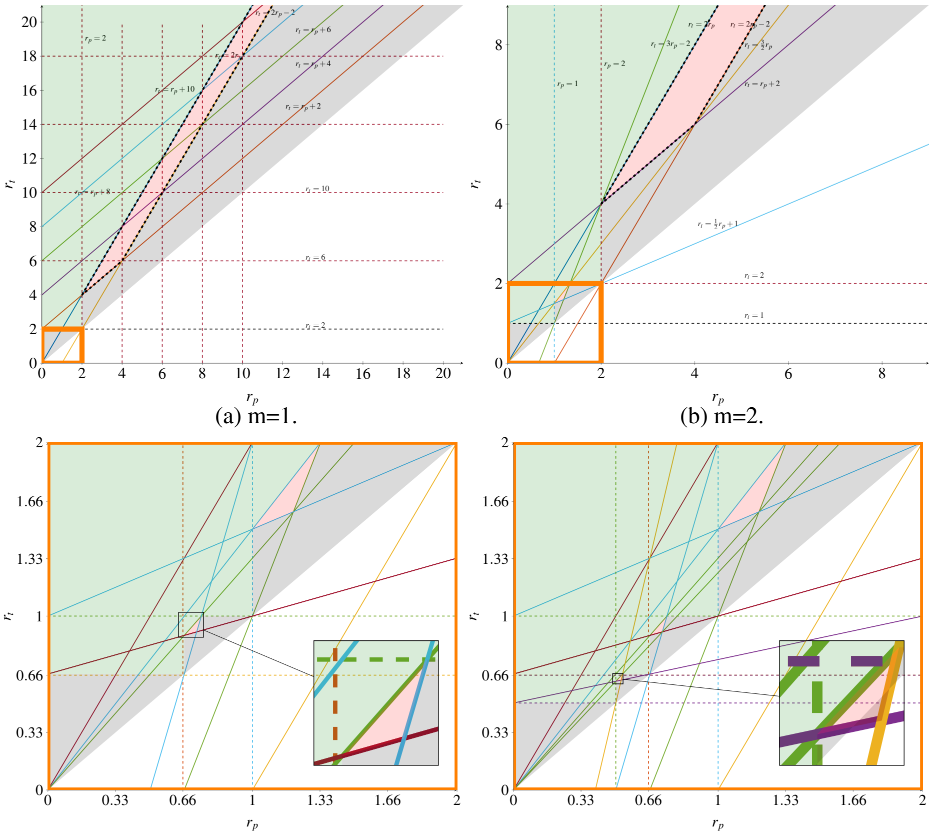

(a) m=1.

(b) m=2.

(c) m=3.

(d) m=4.

Fig. 5: The problem parameter space and the existence of strategies for robots with differing m. The green region depicts **PP** and **TT** conditions, where a suitable strategy exists no matter the initial I-state. The gray region represents the conditions that violate either **PP** or **TT** (or both). The pink region depicts conditions known to be I-state dependent. The white region represents trivially infeasible problems. The orange rectangles emphasize the different magnification-levels and highlight conditions where the differences in sensing power come into play. Both r_p and r_t are expressed in units of $\frac{\delta}{2}$.

Preserving the privacy of a target certainly makes some unusual demands of a sensor. O'Kane's quadrant sensor has pre-images for each output class that are infinite subsets of the plane, making it possible for his robot to increase its uncertainty if it must do so. But it remains far from clear how one might tackle the same problem with a different sensor. The privacy requirement makes it difficult to reason about the relationship between two similar sensors. For example, an octant sensor appears to be at least as useful as the quadrant sensor, but it makes preserving privacy rather trickier. Since octants meet at the origin at 45°, it is difficult to position the robot so that it does not discover too much. An advantage of the one-dimensional model is that the parameter m allows for a natural modification of sensor capabilities. This leads to three closely related results, each of which helps clarify how certain limitations persist even when m is increased.

Theorem 1.A ***(More sensing won't grant omnipotence)*** *The 1-dim. robot is not always able to achieve privacy-preserving tracking, regardless of its sensing power.*

Proof. The negative results in Lemmas 2 and 4 show that there are circumstances where it is impossible to find a tracking strategy satisfying both *PPC* and *TTC*. Though these regions depend on m, no finite value of m causes these regions to be empty. □

Turning cases that depend on the initial I-state,[†] one might think that sensitivity to η_0 is a consequence of uncertainty that compounds because the sensors used are too feeble. But this explanation is actually erroneous. Observe that the I-state dependent region is more complicated than other regions within the strategy space: in Figure 5, the green region is contiguous, whereas the regions marked pink are not. (Bear in mind that the figure does not depict the η_0 itself, merely regions where it is known that η_0 affects the existence of a strategy.) The specific I-state dependent region for Lemma 5 under condition $a = 1$, visible clearly as chisel shape in Figures 5a and 5b, is invariant with respect to m, so remains I-state dependent in all circumstances (though outside the visible region in Figures 5c and 5d, it is present). As m increases, what happens is that the regions formerly marked as not both **PP** and **TT** are claimed as **PP** and **TT**, or become I-state dependent. The following expresses the fact that additional sensing power fails to reduce the number of I-state dependent strategy regions.

Theorem 1.B ***(Information State Invariance)*** *The number of initial I-state dependent strategy regions does not decrease by using more powerful sensors.*

Proof. We focus on the I-state dependent areas within the square between $(0,0)$ and $(2,2)$ as the parts outside this (orange) square do not change as m increases. According to Lemma 5, the I-state dependent conditions for any specific a are bounded by the following linear inequalities:

$$r_t < \frac{2}{a-1}, \qquad (1)$$

$$r_p > \frac{2}{a}, \qquad (2)$$

$$r_t \geq \frac{r_p}{a} + \frac{2}{a}, \qquad (3)$$

$$r_t \geq (a+1)r_p - 2, \qquad (4)$$

$$r_t < \frac{(a+1)}{a} r_p. \qquad (5)$$

[†] For brevity sometimes we will call such instances "I-state dependent" though it is strictly $|\eta_0|$, the size of the initial I-state, on which they depend.

Combining (1)–(3) gives both the bound for r_t as $r_t \in [\frac{2(a+1)}{a^2}, \frac{2}{a-1})$, and the bound for r_p as $r_p \in [\frac{2}{a}, \frac{2a}{a^2-1}]$. The I-state dependent condition, thus, is a bounded region.

Next, we show that (1) and (2) are dominated by (3)–(5). According to (4) and (5), we have $r_p < \frac{2a}{a^2-1}$. Applying this result to (5) produces (1). Similarly, combining (3) and (5) together yields (2). Hence, the I-state dependent conditions are fully determined by inequalities (3)–(5). (Figure 6 provides a visual example.)

To form a bounded region with three linear inequalities, the I-state region has to be a triangle. The three points of the triangle can be obtained by intersecting pairs of (3)–(5): $(\frac{2}{a}, \frac{2a+2}{a^2})$, $(\frac{2a}{a^2-1}, \frac{2}{a-1})$, $(\frac{2(a+1)}{a^2+a-1}, \frac{2(a+1)^2}{a^2+a-1} - 2)$. Since $a \in \{2, 3, \cdots, m\}$, the triangle region will not be empty. Let $\Delta(a)$ denote the triangle with parameter a. Then the smallest y coordinate for $\Delta(a)$ is $minY(\Delta(a)) = \frac{2a+2}{a^2}$. And the largest y coordinate for $\Delta(a)$ is $maxY(\Delta(a)) = \frac{2}{a-1}$. For adjacent triangles $\Delta(a)$ and $\Delta(a+1)$, we have $minY(\Delta(a)) > maxY(\Delta(a+1))$. Hence, the triangles for different values of a do not overlap. □

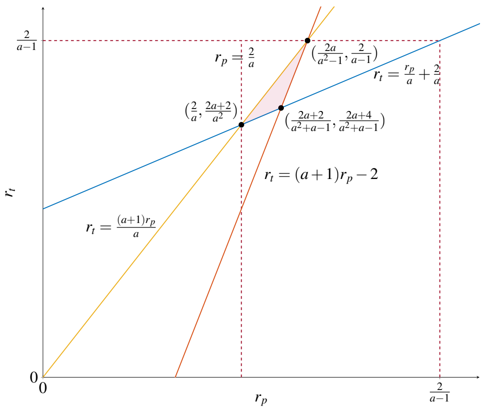
$\frac{2}{a-1}$
$r_p = \frac{2}{a}$
$(\frac{2a}{a^2-1}, \frac{2}{a-1})$
$r_t = \frac{r_p}{a} + \frac{2}{a}$
$(\frac{2}{a}, \frac{2a+2}{a^2})$
$(\frac{2a+2}{a^2+a-1}, \frac{2a+4}{a^2+a-1})$
$r_t = (a+1)r_p - 2$
r_t
$r_t = \frac{(a+1)r_p}{a}$
0
0
r_p
$\frac{2}{a-1}$

Fig. 6: Relationships of the linear inequalities in Lemma 5.

The preceding discussion showed that the number of I-state dependent regions increases with m and that, along with the chisel shaped region, there are $m-1$ triangles of decreasing size. This motivates our introduction of a quantitative measure of the robot's power as a function of m.

Definition 4. *A measure of tracking power, $p(m)$, for the robot with m sensing parameters should satisfy the following two properties: $p(m) > 0$, $\forall m \in \mathbb{Z}^+$* (positivity)*, and $p(a) > p(b)$, if $a, b \in \mathbb{Z}^+$ and $a > b$* (monotonicity).

The plots in Figure 5 suggest that one way to quantify change in these regions is to measure changes in the various areas of the parameter space as m increases. As a specific measure of power in the one-dimensional setting, we might consider the proportion of cases (in the r_p vs. r_t plane) that are **PP** and **TT** (green) and I-state dependent (pink). Though the green and pink areas are unbounded in the full plane, the only changes that occur as m increases are in the square between $(0,0)$ and $(2,2)$. Thus, we take $p(m)$ to equal to the total volume of green and pink regions filling within the region $0 \leq r_p \leq r_t$ and $0 \leq r_t \leq 2$. This area satisfies the properties in Definition 4 and is indicative of the power of the robot as, intuitively, it can be interpreted as an upper-bound of the solvable cases.

Corollary 1 ***(Asymptotic tracking power)*** *The power $p(m)$ of a robot with m sensing parameters to achieve privacy-preserving tracking in the 1-dim. problem is bounded and $\lim_{m\to\infty} p(m) = L$, with $1.5 < L < 1.6$*

Proof. Inequalities (3)–(5) in the proof of Theorem 1.B give $m-1$ triangles, one for each $a \in \{2,3,\cdots,m\}$. An analytic expression gives the area of each of these triangles and the series describing the cumulative pink volume $p_{pink}(m)$ within $0 \leq r_p \leq r_t$ and $0 \leq r_t \leq 2$ can be shown (by the comparison test) to converge as $m \to \infty$. Similarly, the cumulative green volume $p_{green}(m)$ within $0 \leq r_p \leq r_t$ and $0 \leq r_t \leq 2$ converges. Numerical evaluation gives the value of the limit $\approx 1.54\bar{5}$ □

4 Beyond one-dimensional tracking

The inspiration for this work was the 2-dimensional case. This section lifts the impossibility result to higher dimensions.

4.1 Mapping from high dimension to one dimension

In the n-dimensional privacy-preserving tracking problem, the state for the panda becomes a point in $\mathbb{R}^n$. The panda can move with a maximum distance of $\frac{\delta}{2}$ in any direction in $\mathbb{R}^n$ within a single time-step, so that the panda's actions fill an n-dimensional ball. The privacy and tracking bound are also generalized from an interval of size r_p and r_t, to an n-dimensional ball of diameter r_p and r_t respectively. That is, the I-state should contain a ball of diameter r_p and be contained in a ball of diameter r_t, so as to achieve privacy-preserving tracking. The robot inhabits the n-dimensional space as well, and attention must be paid to its orientation too.

It is unclear what would form the appropriate higher dimensional analogue of parameter m, so we only consider n-dimensional tracking problems for robots equipped with a generalization of the quadrant sensor. The sensor's orientation is determined by that of the robot and it indicates which of the 2^n possible orthogonal cells the panda might be in. Adopting notation and definitions analogous to those earlier, we use a tuple for n-dimensional tracking problems—a subscript makes the intended dimensionality clear.

The following lemma shows that there is a mapping which preserves the tracking property from n-dimensional problem to 1-dimensional problems.

Lemma 6. *Given some 1-dim. panda tracking problem $P_1 = (\eta_0, r_p, r_t, n)$, there exists an n-dim. panda tracking problem $P_n = (\theta_0, r_p, r_t, \delta, 1^n)$ where, if $\boldsymbol{TT}(P_n)$, then $\boldsymbol{TT}(P_1)$.*

Proof. The approach to this proof has elements of a strategy stealing argument and simulation of one system by another. The robot faced with a 1-dim. problem constructs an n-dim. problem and uses the (hypothetical) strategy for this latter problem to select actions. The crux of the proof is that the 1-dim. robot can report back observations that are apposite for the n-dim. case. (Figure 7, below, gives a visual overview.)

For some $P_1 = (\eta_0, r_p, r_t, \delta, m)$, with $m = n$, we construct $P_n = (\theta_0, r_p, r_t, \delta, 1^n)$ as follows. Without sacrifice of generality, assume that in P_1 the initial I-state $\eta_0 = \{x | \eta_0^{min} \leq x \leq \eta_0^{max}\}$ is centered at the origin, so $\eta_0^{min} = -\eta_0^{max}$. (This simplifies the argument and a suitable translation of coordinate system rectifies the situation otherwise.) Then we choose θ_0 as the closed ball at the origin with radius η_0^{max}.

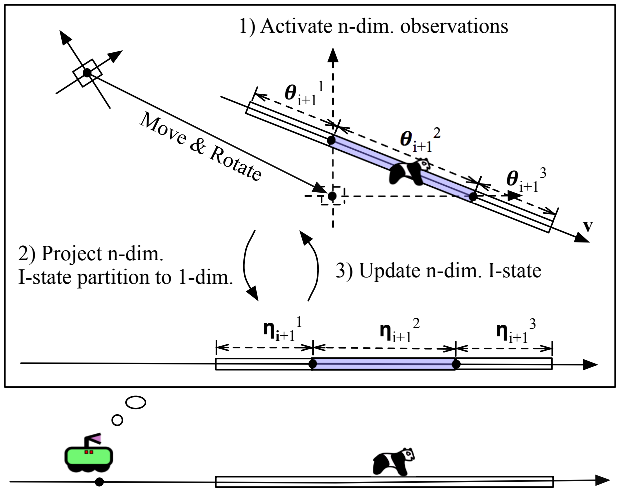

Fig. 7: Constructing a 1-dim. strategy π_1 from some n-dim. strategy π_n.

We show how, given some π_n on $P_n = (\theta_0, r_p, r_t, \delta, 1^n)$, we can use it to define a π_1 for use by the 1-dim. robot. The robot forms θ_0 and also has η_0. It picks an arbitrary unit-length vector $\hat{\mathbf{v}} = v_1\mathbf{e_1} + v_2\mathbf{e_2} + \cdots + v_n\mathbf{e_n}$, unknown to the source of π_n, which is the subspace that the 1-dim. panda lives in. For subsequent steps, the robot maintains $\theta_1^-, \theta_1, \theta_2^-, \theta_2, \ldots, \theta_k^-, \theta_k, \theta_{k+1}^-, \ldots$ along with the I-states in the original 1-dim. problem $\eta_1^-, \eta_1, \eta_2^-, \eta_2, \ldots, \eta_k^-, \eta_k, \eta_{k+1}^- \ldots$. For any step k, the η_k can be seen as measured along $\hat{\mathbf{v}}$ within the higher dimensional space. Given θ_{k-1}, θ_k^- is constructed using Minkowski sum operations as before, though now in higher dimension. Given θ_k^-, strategy π_n determines a new pose for the n-dim. robot and, on the basis of this location and orientation, the n sensing planes slice through θ_k^-. Though the planes demarcate 2^n cells, the line along $\hat{\mathbf{v}}$ is cut into no more than $n+1$ pieces as the line can pierce each plane at most once (with any planes containing $\hat{\mathbf{v}}$ being ignored). Since the 1-dim. robot has $m = n$, it picks the $u_1, \ldots, u_n$ by measuring the locations that the sensing planes intersect the line $\mathbf{x} = \alpha\hat{\mathbf{v}}$, $\alpha \in \mathbb{R}$. (If fewer than n intersections occur, owing to planes containing the line, the extra u_i's are simply placed outside the range of the

I-state.) After the 1-dim. panda's location is determined, the appropriate orthogonal cell is reported as the n-dim. observation, and θ_k^- leads to θ_k via the intersection operation. This process comprises π_1. It continues indefinitely because θ_k must always entirely contain η_k along the line through $\hat{\mathbf{v}}$ because, after all, a cantankerous n-dim. panda is free to choose to always limit its movements to that line.

If π_n is **TT**, then so too is the resulting strategy π_1 since the transformation relating $\theta_k \cap \{\alpha\hat{\mathbf{v}} : \alpha \in \mathbb{R}\}$ with η_k preserves length and, thus, θ_k fitting within a ball of diameter r_t implies that $|\eta_k| < r_t$. □

4.2 Impossibility in high-dimensional privacy-preserving tracking

Now we are ready to connect the pieces together for the main result:

Theorem 2 ***(Impossibility)*** *It is not possible to achieve privacy-preserving panda tracking in n dimensions for every problem with $r_p < r_t$.*

Proof. To extend the lemmas that have shown this result for $n = 1$ to cases for $n > 1$, suppose such a solution existed for $P_n = (\theta_0, r_p, r_t, \delta, 1^n)$. Then according to Lemma 6, every 1-dim. panda tracking problem $P_1 = (\eta_0, r_p, r_t, \delta, n)$ is **TT**, since they can be mapped to an n-dim. **TT** panda tracking problem. But this contradicts the non-**TT** instances in Lemma 2, so no such strategy can exist for every non-trivial problem ($r_p < r_t$) in two, three, and higher dimensions. □

5 Related work

Information processing is a critical part of what most robots do—without estimating properties of the world, their usefulness is hampered, sometimes severely so. Granting our computational devices access to too much information involves some risk, and privacy conscious users may balk and simply opt to forgo the technology. Multiple models have been proposed to think about this tension in setting of data processing more generally, *cf.* formulations in [5,6]. Existing work has also explored privacy for networks, along with routing protocols proposed to increase anonymity [7], and includes wireless sensor problems where distinct notions of spatial and temporal privacy [8,9] have been identified.

Some recent work has begun to investigate privacy in settings more directly applicable to robots. Specifically three lines of work propose techniques to help automate the development of controllers that respect some limits on the information they divulge. O'Kane and Shell [10] formulated a version of the problem in the setting of combinatorial filters where the designer provides additional stipulations describing which pieces of information should always be plausibly indistinguishable and which must never be conflated. The paper describes hardness results and an algorithm for ascertaining whether a design satisfying such a stipulation is feasible. This determination of feasibility for a given set of privacy and utility choices is close in spirit to what this paper has explored for the particular case of privacy-preserving tracking: here those choices become quantities r_p and r_t. The second important line is that of Wu et al. who explore how to disguise some behavior in their control system's plant model, and show how to protect this secret behaviour by obfuscating the system outputs [11]. More recent work expresses both utility and privacy requirements in an automata framework and proposes algorithms to synthesize controllers satisfying these two constraints [12]. In the third line, Prorok and Kumar [13] adopt the differential privacy model to characterize the privacy of heterogeneous robot swarms, so that any individual robot cannot be determined to be of a particular type from macroscopic observations of the swarm.

Finally, we note that while we have considered panda tracking as a cute realization of this broader informationally constrained problem, wildlife monitoring and protection is an area in which serious prior robotics work exists (e.g., see [14]) and for which there is substantial and growing interest [15].

6 Conclusion and future work

In this paper we have reexamined the panda tracking scenario introduced by O'Kane [3], focusing on how various parameters the specify a problem instance (such as the capabilities of the robot and the panda) affect the existence of solutions. Our approach has been to study nontrivial instances of the problem in one dimension. This allows for an analysis of strategies by examining whether the sensing operations involved at each step increase or decrease the degree of uncertainty in a directly quantifiable way. Only if this uncertainty can be precisely controlled forever, can we deem the problem instance solved. We use a particular set of sensing choices, basically division of the region evenly, which we think of as a *split* operation. This operation is useful because the worst resulting I-state in other choices, those that are not evenly split, is weaker (in terms of satisfying the tracking and privacy constraints) than even divisions are. Thus, the split operation acts as a kind of basis: if the problem has no solution with these choices, the panda cannot be tracked with other choices either.

In examining the space of tracking and privacy stipulations, the existence of strategies is shown to be a function of the robot's initial belief and panda's movement. There exist regions without any solution, where it is impossible for the robot to actively track the panda as well as protect its privacy for certain nontrivial tracking and privacy bounds. Additionally, we have uncovered regions where solution feasibility is sensitive to the robot's initial belief, which we have called I-state dependent cases (or conditions). The simple one-dimensional setting also permits exploration of how circumstances change as the robot's sensing power increases. Perhaps surprisingly, the number of these I-state dependent strategy conditions does not decrease as the robot's sensing becomes more powerful. Finally, we connect the impossibility result back to O'Kane's setting by mapping between high-dimensional and one-dimensional versions, proving that the 2D planar panda tracking problem does not have any privacy-preserving tracking strategy for every non-trivial tracking and privacy stipulation.

The results presented in this paper reveal some properties of a particular—and it must be said somewhat narrow—instance of an informationally constrained problem. Future work could explore what sensor properties permit such a task to be achieved, perhaps identifying families of sensors that suffice more broadly. Also, probabilistic models impose different belief representations and observation functions and it is worth exploring analogous notions in those settings. Another thread is to weaken the poacher's capability, allowing the privacy bound to be relaxed somewhat. For example, in considering some latency in the poacher's reaction to information that has been leaked, it may be acceptable for the privacy bound to be relaxed so that it need not apply from every time-step to the next.

Acknowledgements

This work was supported by the NSF through awards IIS-1527436 and IIS-1453652. We thank the anonymous reviewers for their time and extraordinarily valuable comments, and are grateful to the conference committee for a student travel award.

References

1. B. R. Donald, "On information invariants in robotics," *Artificial Intelligence — Special Volume on Computational Research on Interaction and Agency, Part 1*, vol. 72, no. 1–2, pp. 217–304, 1995.
2. M. T. Mason, "Kicking the sensing habit," *AI Magazine*, vol. 14, no. 1, pp. 58–59, 1993.
3. J. M. O'Kane, "On the value of ignorance: Balancing tracking and privacy using a two-bit sensor," in *Proc. Int. Workshop on the Algorithmic Foundations of Robotics (WAFR'08)*, Guanajuato, Mexico, 2008.
4. S. M. LaValle, "Sensing and filtering: A fresh perspective based on preimages and information spaces," *Foundations and Trends in Robotics*, vol. 1, no. 4, pp. 253–372, 2010.
5. L. Sweeney, "*k*-anonymity: A model for protecting privacy," *Int. Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, vol. 10, no. 5, pp. 557–570, 2002.
6. C. Dwork, "Differential privacy: A survey of results," in *Proc. Int. Conference on Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.
7. I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong, "Freenet: A distributed anonymous information storage and retrieval system," in *Int. Workshop on Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Unobservability*, 2001, pp. 46–66.
8. P. Kamat, Y. Zhang, W. Trappe, and C. Ozturk, "Enhancing source-location privacy in sensor network routing," in *Proc. IEEE Int. Conference on Distributed Computing Systems (ICDCS'05)*, Columbus, Ohio, USA, 2005, pp. 599–608.
9. P. Kamat, W. Xu, W. Trappe, and Y. Zhang, "Temporal privacy in wireless sensor networks," in *Proc. Int. Conference on Distributed Computing Systems (ICDCS'07)*, Toronto, Ontario, Canada, 2007, pp. 23–23.
10. J. O'Kane and D. Shell, "Automatic design of discreet discrete filters," in *Proc. IEEE Int. Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, 2015, pp. 353–360.
11. Y.-C. Wu and S. Lafortune, "Synthesis of insertion functions for enforcement of opacity security properties," *Automatica*, vol. 50, no. 5, pp. 1336–1348, 2014.
12. Y.-C. Wu, V. Raman, S. Lafortune, and S. A. Seshia, "Obfuscator synthesis for privacy and utility," in *NASA Formal Methods Symposium*. Springer, 2016, pp. 133–149.
13. A. Prorok and V. Kumar, "A macroscopic privacy model for heterogeneous robot swarms," in *Proc. Int. Conference on Swarm Intelligence (ICSI)*. Springer, 2016, pp. 15–27.
14. D. Bhadauria, V. Isler, A. Studenski, and P. Tokekar, "A robotic sensor network for monitoring carp in minnesota lakes," in *Proc. IEEE Int. Conference on Robotics and Automation (ICRA)*, Anchorage, AK, USA, 2010, pp. 3837–3842.
15. F. Fang, P. Stone, and M. Tambe, "When security games go green: Designing defender strategies to prevent poaching and illegal fishing," in *Proc. Int. Joint Conference on Artificial Intelligence (IJCAI)*, Buenos Aires, Argentina, 2015.