Indistinguishability in Localization and Control with Coarse Information

Daniel Liberzon Sayan Mitra
University of Illinois Urbana-Champaign
Coordinated Science Laboratory
Urbana, IL, USA
{liberzon,mitras}@illinois.edu

Abstract

We study localization and control problems in which agent dynamics are described by difference or differential equations, while output measurements are collected at discrete times and given by finite-valued maps depending on possibly unknown landmark locations. Guided by the goal of understanding fundamental limitations imposed by such coarse measurements, we focus on characterizing indistinguishable states, i.e., agent-landmark pairs that produce identical observations under all control inputs. We show that indistinguishability relations can be checked automatically under mild assumptions and, being a special type of bisimulation, we develop an iterative algorithm for approximately computing them. We then introduce an analytical approach, rooted in observability theory of linear control systems, which iteratively computes a sequence of subspaces converging in finitely many steps to the indistinguishable subspace; a differential-geometric extension to nonlinear systems is also outlined.

CCS Concepts

• **Theory of computation** → *Timed and hybrid models*; • **Computer systems organization** → **Robotics**.

Keywords

SLAM, observablity, bisimulation, indistinguishable sets

ACM Reference Format:
Daniel Liberzon Sayan Mitra. 2025. Indistinguishability in Localization and Control with Coarse Information. In *28th ACM International Conference on Hybrid Systems: Computation and Control (HSCC '25), May 6–9, 2025, Irvine, CA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3716863.3718034

1 Introduction

The nature of available measurements limit the degree of control that can be achieved in a system. In the traditional analog domain, control performance is studied using smooth measurement models. More recent developments (surveyed, e.g., in [22]) give bounds on the necessary bandwidth and sensor resolutions for quantized and networked control systems using digital communication. With control systems using machine learning-enabled perception of features in complex environments—like edges, keypoints, and landmarks—we aim to investigate the limits of control and estimation with such *coarse measurements*.

Simultaneous Localization and Mapping (SLAM) [8, 28] is a prominent task that has to be solved with coarse measurements: an agent has to build a map of an unknown environment based on available features, while simultaneously determining its position in that map. From earlier Extended Kalman Filter (EKF)-based approaches, modern SLAM algorithms have evolved to use advanced methods like particle filters, graph optimization, and direct sparse odometry [9], and have become integral components in robotic vacuums, Mars rovers [5], and self-driving cars [31]. While SLAM implementations differ, there are universal limitations to how far these algorithms can reduce uncertainty in ambiguous or feature-poor environments. One well-known situation is where the agent has to distinguish two locations in a symmetric hallway that produce similar measurements. Another version of this is the well-known *loop closure problem* in which the agent has to detect that it has returned to a previously visited location [30].

In this paper, we examine the fundamental limits of algorithms solving a simplified version of SLAM with a single landmark, through the lens of nondeterministic or set-valued uncertainty. While there is extensive empirical research on SLAM, the hybrid systems point of view can bring a fresh perspective since it is characterized by interactions between the agent's continuous state space and dynamics, discrete or quantized measurements, and the need to effectively reason about uncertainty propagation and quantification.

To fix ideas, consider a map with a set of landmarks in unknown locations, and an agent with a binary sensor that gives a positive measurement if it is close to any landmarks. A typical SLAM algorithm works via alternating predictions and corrections: Based on the current uncertainty S_t (in agent's position and the position of the landmarks), a control action u_t is chosen, and then, based on the actual measurement y_t received after applying u_t the estimate is corrected to S_{t+1}. Given some initial uncertainty S_0, we want to know what is the best *any* SLAM algorithm can do in reducing the uncertainty?

For an algorithm to reduce the uncertainty S_t (jointly in position of the agent and the landmarks), it has to choose the control action u_t that yields a measurement y_t that eliminates some state (x, m)—agent and landmark position combination—from S_t as a possibility or a counter-factual. This observation leads to the notion of *distinguishability* of states: two states (x_1, m_1) and (x_2, m_2) are distinguishable if there exists a sequence of control inputs $\overline{u}$

This research was supported by AFOSR MURI grant (FA9550-23-1-0337) entitled HyDDRA: Hybrid Dynamics - Deconstruction and Aggregation.

HSCC '25, Irvine, CA, USA

ACM ISBN 979-8-4007-1504-4/2025/05
https://doi.org/10.1145/3716863.3718034

which, when applied to each, leads to different observation sequences. In contrast, states are *indistinguishable* if no such control sequence exists. Thus, understanding the limits of any algorithm tantamounts to characterizing the indistinguishable states. We pursue both a computational and an analytic approach towards this characterization.

After introducing indistinguishability formally, we show that for a general class of dynamic and measurement functions, putative indistinguishability relations can be verified automatically by solving satisfiability problems. We also observe that indistinguishability relations are a special type of bisimulation relations [11, 21, 26], and provide an iterative algorithm for approximating them.

Our analytic approach is rooted in Kalman's observability decomposition for linear systems (see, e.g., [12]) and its differential-geometric extension to nonlinear systems [13]. In this approach, a sequence of subspaces (or, more generally, (co)distributions) is iteratively computed and converges in finitely many steps to yield a closed-form description of the indistinguishable set. Prior work on using system-theoretic notion of observability for SLAM [1, 17] do not consider uncertainty in the landmarks, just in the position of the robot. These works establish the lack of observability by standard computations, but do not analyze the unobservable subspaces, instead their focus is on finding the sufficient number of measurements for gaining full observability.

Our work is inspired by [30] in which indistinguishability of robot environments has been characterized from the point of view of continuity of SLAM and filtering algorithms as maps on information spaces [16, 29]. However, our work provides computational and analytical techniques for computing indistinguishable sets of state-landmark pairs for a specific class of SLAM scenarios, in contrast to the more general and more abstract topological characterizations of indistinguishable environments given in [30]. Our focus on coarse measurements also distinguishes this paper from earlier work such as [17], where output is not quantized, and puts it more in line with the recent literature on estimation and control with minimal data (see, e.g., [18, 22, 24] and the references therein).

Research on control with partial observations utilizes similar concepts as our work in addressing a different problem. For example, in [15], the authors introduce a detector akin to a state estimator which is used by the downstream controller to meet the desired specifications. In [20], the authors construct a controller using the idea of a knowledge abstraction. These works are concerned with specifications defined in terms of finite or omega-regular sequences of observations, and it is unclear whether our objective of reducing the uncertainty in the state estimation can be written as such a specification. The notion of detectability in [15] is related to indistinguishability, in that if a system is detectable then the indistinguishable states are trivial singletons.

In summary, the main contributions of this paper are as follows. First, we characterize indistinguishable sets of a simple class of SLAM problems in Propositions 3.7 and 3.8. We give general methods for checking and computing indistinguishability relations via SMT solving in Proposition 4.1, a bisimulation-type semi-algorithm in Proposition 4.2, and through Kalman observability decomposition in Proposition 5.1. Finally, in Proposition 3.6 we relate indistinguishable sets to quality of solutions achievable by SLAM.

2 Example: Localization using a landmark

Consider an agent described by a point $x(t)$ moving counterclockwise on a circle centered at 0 in the plane, according to the dynamics

$$\dot{x}_1 = -x_2, \quad \dot{x}_2 = x_1 \tag{1}$$

A *landmark* will be a fixed point m in the plane, different from 0. We take quantized and sampled measurements of the distance from $x(t)$ to the landmark. Geometrically, this means constructing a finite number of concentric circles around m and being told which annulus contains $x(t)$. Here we consider the simplest—but also the most practical—case where there is only one such circle, and we know whether $x(t)$ is inside or outside. (This corresponds to trying to establish a signal connection with the landmark, and returning 1 if the landmark is in range and 0 otherwise.) We denote by S_m this circle around m (whose radius may correspond to signal range) and by S_x the circle containing $x(t)$. We assume that the two circles S_x and S_m are *both known* (at least for the moment) and intersect in a nontrivial way. We take the measurements every T units of time, where T is the sampling period which we take to be a rational number.

Note that in terms of data rate, this means that the agent is sending[1] 1 bit every T time units. Increasing T, we can make the data rate as small as we want. Since the system (1) is (marginally) stable and hence has zero entropy, we indeed expect to be able to asymptotically localize $x(t)$ with arbitrarily small data rate (see, e.g., [18] and the references therein).

At $T = 0$, the measurement divides the circle S_x into two arcs—the one inside S_m and its complement—and tells us which of these two arcs contains $x(0)$. This arc represents the initial uncertainty set. Using the dynamics (1), it is easy to propagate this arc forward, i.e., generate reachable sets from this arc for future times. At each subsequent sampling time T, $2T$, ... we have a possibility of reducing the uncertainty if the reachable set at that time nontrivially intersects S_m.

Suppose that at some time t we have localized $x(t)$ to within an arc of length ε, centered at some point which we label as $\hat{x}(t)$. Considering the trajectory of (1) passing through $\hat{x}(t)$ at time t and sampling it at the subsequent sampling times $kT > t$ generates a discrete orbit on the circle S_x. Since the period of rotation for x is 2π and we took T to be rational, the two numbers are rationally independent in the sense of the condition (20) in the Appendix. Thus Lemma 1 from the Appendix applies, telling us that this orbit is dense everywhere on the circle. This implies that there will be a sampling time, say k^*T, at which the corresponding point on the orbit will be within, say, $\varepsilon/4$ of a point of intersection between S_x and S_m. At that time, it is easy to see that the measurement will reduce the uncertainty set to an arc of length at most $3\varepsilon/4$ (in the worst case). Repeating this argument, we see that the localization error will asymptotically converge to 0.

The previous discussion assumes that we have not been acting on the measurements that have come between time t and time k^*T, i.e., no correction updates to the uncertainty set have been executed during that time period. This is not realistic, as in practice

[1] Depending on the specific set-up, we can think of the agent as using a communication channel to transmit encoded messages to a remote base or a central decoder/controller unit.

we don't know when the above condition defining the time k^*T is fulfilled, instead, we simply make an update whenever we receive a signal from the landmark. However, this does not invalidate the above conclusion. To see why, consider two uncertainty sets: the actual one (updated at every sampling time) and the hypothetical enlarged one, defined (for analysis purposes only) by ignoring the updates until the time k^*T defined via the condition above. Clearly, the actual uncertainty set is always contained in the enlarged one. We showed that the size of the enlarged uncertainty set will converge to 0. Hence, so will the size of the actual uncertainty set. Our conclusions for this example are as follows:

(1) Different initial conditions produce different sequences of output measurements.
(2) For each initial condition $x(0)$, the estimate $\hat{x}(\cdot)$ produced by the above procedure asymptotically converges to the corresponding state trajectory $x(\cdot)$, i.e., $\hat{x}(t) - x(t) \to 0$ as $t \to \infty$.

Note that property (1) above, while not explicitly established in the preceding argument, can be derived from property (2) because the procedure relies on output measurements only. This situation is in contrast with scenarios that we will encounter in the rest of the paper, where different initial conditions may produce identical output sequences and are thus *indistinguishable* based on output measurements (and as a consequence, property (2) also cannot be achieved).

3 SLAM and Indistinguishability

For a set $\mathcal{X}$, we denote its k-fold Cartesian product as $\mathcal{X}^k$. Given a relation $R \subseteq \mathcal{X} \times \mathcal{X}$ on $\mathcal{X}$ and an element $x \in \mathcal{X}$, we denote by $R(x) = \{x' \in \mathcal{X} \mid (x, x') \in R\}$.

Consider a discrete-time system described by the equations:

$$x_{t+1} = f(x_t, u_t) \qquad y_t = h(x_t). \tag{2}$$

Here $x_t \in \mathcal{X}$ is the state of the system at time t, $\mathcal{X} \subseteq \mathbb{R}^n$ is the state space, $u_t \in \mathcal{U}$ is the control action applied at t, and y_t is a *measurement* (or an observation) made in the state x_t. The point of departure in our work from the standard setup is that h is a nonsmooth and possibly finite-valued *measurement function*, which makes (2) relevant from the point of view of hybrid systems.

Given an initial state x_0, an infinite sequence of control actions $\overline{u} = u_0, u_1, \ldots$, generates a corresponding sequence of states $x_0^{\overline{u}} = x_0, x_1, \ldots$, and a sequence of measurements $h(x_0^{\overline{u}}) = y_0, y_1, \ldots$. The state x_0 is the initial state and y_0 is the initial measurement, prior to the application of any control. The k^{th} elements in these sequences are denoted by $\overline{u}[k], x_0^{\overline{u}}[k]$, and $x_0^{\overline{u}}[k]$.

3.1 A simplified SLAM problem

Consider a robot with a sensor moving in an Euclidean space with k stationary landmarks $m^1, \ldots, m^k \in \mathbb{R}^n$. At every step, the robot gets a k-bit measurement in which the i^{th}-bit is 1 if and only if the robot's position x is within r^i-distance of m^i. That is,

$$\begin{aligned} x_{t+1} &= f(x_t, u_t) \\ m_{t+1}^i &= m_t^i, \forall\, i = 1, \ldots, k, \\ y_t &= h(x_t, m_t), \end{aligned} \tag{3}$$

where $h^i(x, m) = 1$ if $|x - m^i| \leq r^i$, and 0 otherwise. Here $m_t = \langle m_t^1, \ldots, m_t^k \rangle$ is the stationary vector of the k landmarks, and r^i is the radius associated with the i^{th} landmark. We can view the stacked vector (x, m) as the full state of the system as defined in (2). Sometimes it will be convenient to alternatively express the state space $\mathcal{X}$ as the Cartesian product $\mathcal{X} \times \mathcal{M}^k$ of the spaces occupied by the robot and the landmarks. The convention being used should be clear from context. This simplified measurement function h gives a *vector* of measurements—one for each landmark. With a more realistic measurement function that only provides the cardinality of the set $\{i \mid h^i(x, m) = 1\}$ of positive measurements, the SLAM algorithms would need address the *data association* problem, which is beyond the scope of the current paper.

Example 3.1 (nDSLAM). We will use instances of the simplified SLAM problem with a single landmark to illustrate some of the results. In nDSLAM, $\mathcal{X}_0 \subseteq \mathbb{R}^n$ and there is a single landmark $m \in \mathcal{M}_0 \subseteq \mathbb{R}^n$. Further, the dynamics is $f(x_t, u_t) = x_t + u_t$, where u_t is an n-dimensional control vector. For $n = 1$ we refer to this problem as 1DSLAM.

In this nonprobabilistic setting, a SLAM algorithm iteratively computes a sequence of control actions such that given some prior initial estimates $\mathcal{X}_0 \subseteq \mathcal{X}, \mathcal{M}_0^i \subseteq \mathcal{M}$ about the unknown position of the robot x_0 and the (static) landmarks m^i, such that, $x_0 \in \mathcal{X}_0$ and $m^i \in \mathcal{M}_0^i$, the uncertainty in the position of the robot and the landmarks is simultaneously reduced, based on the sequence of measurements.

We formally define a generic SLAM algorithm as a function g of the following type, whose computation is prepended to (2) in completing the system description:

$$u_t, S_{t+1}, v_{t+1} = g(y_t, S_t, v_t). \tag{4}$$

Here $S_0 = \mathcal{X}_0$ (or $\mathcal{X}_0 \times \mathcal{M}_0$) is the initial joint uncertainty in the position of the robot and the landmarks, v_0 is the initial valuation of internal program variables, and y_0, as before, is the initial measurement from the state x_0. For each time step t, the algorithm g updates the joint uncertainty to S_{t+1}, the internal variables to v_{t+1} for the next time step, and it computes the control input u_t for the current time step, which is then fed to (2) to produce the state x_{t+1} and the measurement y_{t+1}. For a given set of initial conditions $(x_0, \mathcal{X}_0, S_0, v_0)$, continuing this sequence of updates completely defines the behavior of the system and the resulting sequences of states, measurements, control actions, and estimates.

We say that a SLAM algorithm g is *sound* if at every step, the estimate contains the actual state, i.e., for all t, $x_t \in S_t$.

3.2 Indistinguishability

In this paper we study the concept of indistinguishability as an obstruction to solving the SLAM problem. Related notions of indistinguishability have been used in control theory [2, 13], in proving impossibility results in distributed computing [19], and featured at least as far back as the Leibniz-Clarke correspondences about principles of construction of physical theories [4, 6].

Definition 3.2. An *indistinguishable set* $I \subseteq \mathcal{X}$ is a set of states that produce identical outputs for any control input. That is, for any

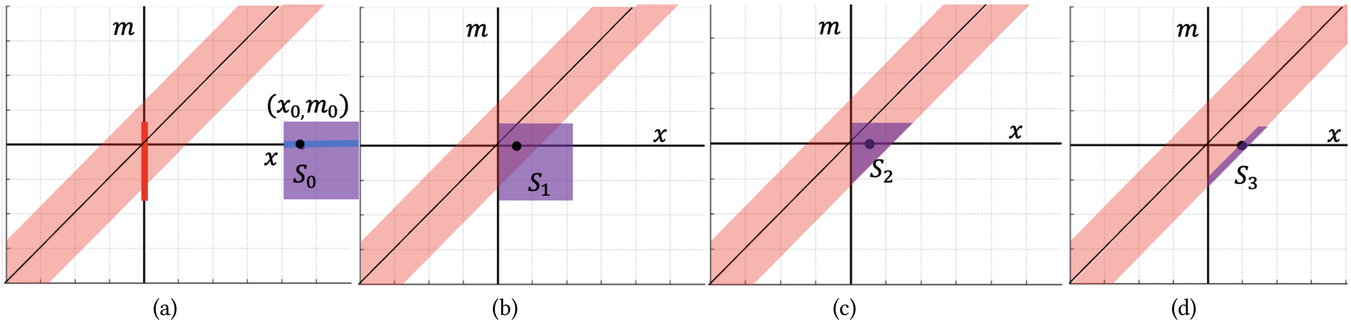

Figure 1: Progression of 1DSLAM algorithm showing agent-landmark positions (x_0, m_0) and the joint uncertainty S_0. (a) The unknown initial position of the agent x_0 is in the known interval $\mathcal{X}_0$ (blue); unknown landmark location m_0 is in the known interval $\mathcal{M}_0$ (red). (b) The SLAM algorithm applies control u_0 which moves the agent to x_1 and the known joint uncertainty to S_1 (purple). (c) The red strip is the region where the output $y = h(x, m) = 1$ and having received the signal $y = 1$ in this execution the SLAM algorithm reduces joint uncertainty to S_2. (d) A SLAM algorithm can move the agent back and forth to reduce the joint uncertainty down to a line segment—an indistinguishable set—but not to a point. This illustrates the limits of solving the SLAM problem in 1D with a single landmark.

$x_1, x_2 \in I$, and any control sequence $\overline{u}$ and any $k \in \mathbb{N}$, $h(x_1^{\overline{u}})[k] = h(x_2^{\overline{u}})[k]$.

In this deterministic setting, a pair of states are *distinguishable* if they are not indistinguishable, i.e., there exists a finite sequence of control inputs applying which they produce different outputs. We say that a pair of states x_1, x_2 are *k-distinguishable* if there exists $\overline{u} \in \mathcal{U}^k$, such that $h(x_1^{\overline{u}})[k] \neq h(x_2^{\overline{u}})[k]$. We can also look at indistinguishability as a relationship on states.

Definition 3.3. An *indistinguishability relation* $\mathcal{I} \subseteq \mathcal{X} \times \mathcal{X}$ is relation that is closed under any control input and such that the related states produce identical observations. That is, for any $(x_1, x_2) \in \mathcal{I}$, (1) $h(x_1) = h(x_2)$ and (2) for any control input u, the corresponding next states satisfy $(f(x_1, u), f(x_2, u)) \in \mathcal{I}$.

Indistinguishability relations are equivalence relations on the states and they are a subset of the set of bisimulation relations [21, 23, 26]. Bisimulations are more permissive in that two states are allowed to preserve observational equivalence using different controls, but for indistinguishability they have to use the *same* controls.

Definition 3.4. A relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{X}$ is a *bisimulation relation* if for any $(x_1, x_2) \in \mathcal{R}$, (1) $h(x_1) = h(x_2)$ and (2) for any control input u_1, there exists a corresponding control input u_2 such that $(f(x_1, u_1), f(x_2, u_2)) \in \mathcal{R}$.

Thus, as we shall see in Section 4, algorithms for computing bisimulation relations can be adapted for computing indistinguishable sets. The complement of an indistinguishability relation is a distinguishability relation $\mathcal{D}$, i.e., if $(x_1, x_2) \in \mathcal{D}$ then there exists an input that leads the two states to produce different outputs.

Proposition 3.5. *For any indistinguishability relation $\mathcal{I}$ and a state* (anchor) $x_0 \in \mathcal{X}$, *$\mathcal{I}(x_0)$ is an indistinguishable set.*

Proof. Consider any state x in the set $\mathcal{I}(x_0)$ and fix any sequence of control inputs $u = u_0, u_1, \ldots$. Now, since $(x_0, x) \in \mathcal{I}$, by condition (1) $h(x_0) = h(x)$ and by condition (2) $(f(x_0, u_0), f(x, u_0)) \in \mathcal{I}$. Now, we can define $x_0' = f(x_0, u_0), x' = f(x, u_0)$, and $u = u_1, u_2, \ldots$, and repeat the above argument for every element in u. This completes the proof. □

The equivalence classes of $\mathcal{I}$ are the indistinguishable sets, i.e., the quotient $\mathcal{X}/\mathcal{I}$ is the set of indistinguishable sets. In other words, an indistinguishability relation $\mathcal{I}$ captures a family of indistinguishable sets, each of which can be identified by different representative states.

We relate the output of SLAM algorithms and indistinguishability as follows. Let $S_0, \ldots, S_T$ be the sequence of estimates computed by a SLAM algorithm g as per Equation (4) and let $\overline{u}$ be the corresponding computed control sequence. We are interested in the initial position of the agent, and therefore, from S_T calculate the estimate x_0 as the iterative preimage of S_T under f with $\overline{u}$, i.e., $S_{init} := f^{-1}(S_T, \overline{u})$.

Proposition 3.6. *If g is sound then $\mathcal{I}(S_{init}) \subseteq S_{init}$.*

Proof. Let $x_0 \in \mathcal{X}_0$ be the actual unknown initial state. By soundness of g, for all $t \leq T$, $x_0^{\overline{u}}[t] \in S_t$. Suppose, for the sake of contradiction, there exists $x_0' \in \mathcal{I}(S_{init}) \setminus S_{init}$. By the definition of the preimage, $x_0 \in S_{init}$. Since, x_0 and x_0' are indistinguishable, the observations from x_0 and x_0' would be identical, i.e., $h(x_0^{\overline{u}})[t] = h(x_0'^{\overline{u}})[t]$ for all t. Let b be the final step where g computes S_b such that $x_0' \notin f^{-1}(S_b, \overline{u})$, i.e., $x_0'^{\overline{u}}[b] \notin S_b$. Since g only depends on the observations, we can swap x_0' and x_0 in this argument with the result that $x_0^{\overline{u}}[b] \notin S_b$. This contradicts the soundness of g. □

3.3 Indistinguishable sets for nDSLAM

In the nDSLAM problem of Example 3.1, we consider the position of the agent and the landmark lumped together $(x, m) \in \mathcal{X} \times \mathcal{M} = \mathbb{R}^{2n}$ as the state of the system. We show that for any $a \in \mathbb{R}^n$, the affine set of states $S_a = \{(x, m) \mid x - m = a\}$ is an indistinguishable set.

Proposition 3.7. *For any $a \in \mathbb{R}^n$, S_a is an indistinguishable set for nDSLAM.*

Proof. Fix $a \in \mathbb{R}^n$ and consider $(x_1, m_1), (x_2, m_2) \in S_a$. If $h(x_1, m_1) = 1$ then $|x_1 - m_1| \leq r$. Since $(x_1, m_1) \in S_a$, $x_1 - m_1 = a$ and it follows that $|a| \leq r$. Since $(x_2, m_2) \in S_a$ it follows that $|x_2 - m_2| = |a| \leq r$, and therefore, $h(x_2, m_2) = 1$. If $h(x_1, m_1) = 0$, then a similar argument shows that $h(x_2, m_2) = h(x_1, m_1)$. Next we have to show that for any control action $u \in \mathbb{R}^n$, (x_1, m_1) and (x_2, m_2) continue to produce identical outputs.

$$\begin{aligned}
& x_1 - m_1 = x_2 - m_2 && \text{[Both equal } a\text{]} \\
\Rightarrow\ & x_1 + u - m_1 = x_2 + u - m_2 \\
\Rightarrow\ & f(x_1, u) - m_1 = f(x_2, u) - m_2 && \text{[Definition of } f\text{]} \\
\Rightarrow\ & h(f(x_1, u), m_1) = h(f(x_2, u), m_2) && \text{[Definition of } h\text{]}.
\end{aligned}$$

This shows that $(f(x_1, u), m_1)$ and $(f(x_2, u), m_2)$ are produce the same output and both are in S_{a+u}. Since this argument was carried out for arbitrary $(x_1, m_1), (x_2, m_2)$ and u, it follows that S_a is an indistinguishable set. □

Thus, every affine set S_a (for $n = 1$ these are diagonals with slope 1) is an indistinguishable set for the nDSLAM, and any subset $S \subseteq S_a$ is also indistinguishable by the same argument. It is straightforward to check that

$$\mathcal{I} = \{((x_1, m_1), (x_2, m_2)) \mid x_1 - m_1 = x_2 - m_2\} \tag{5}$$

is an indistinguishability relation, and for a particular (x_0, m_0), $\mathcal{I}(x_0, m_0) = S_{x_0 - m_0}$ is the set of states that are indistinguishable from it. The infinite collection of equivalence classes of $\mathcal{I}$ are the diagonal indistinguishable sets S_a, $a \in \mathbb{R}^n$. From the preceding discussion, $\mathcal{I}^c = \{((x_1, m_1), (x_2, m_2)) \mid x_1 - m_1 \neq x_2 - m_2\}$ is a distinguishability relation. This can also be seen in terms of distinguishable sets as follows.

Proposition 3.8. *Any $(x_1, m_1) \in S_a$ and $(x_2, m_2) \notin S_a$ are distinguishable.*

Proof. For $x_i, m_i \in \mathbb{R}^n$, $i = 1, 2$, if $h(x_1, m_1) \neq h(x_2, m_2)$ then the two points are distinguishable and we are done. For the rest of the proof, we assume w.l.o.g., that $h(x_1, m_1) = h(x_2, m_2) = 0$. Let us define $z_i = x_i - m_i$. Let $s = |z_2 - z_1|$. Since $z_1 \neq z_2$, it follows that $s > 0$. Next we define the control input vector u that distinguishes x_1 and x_2 as

$$u = -z_1 + \frac{r}{s}(z_2 - z_1).$$

This u shifts z_1 toward z_2 by a fraction $\frac{r}{s}$ of the vector $z_2 - z_1$. Next we show that with this control input $z_1 + u$ gets an observation of 1 while $z_2 + u$ stays 0. Applying u to (x_1, m_1), we have $f((x_1, m_1), u) = (x_1 + u, m_1)$ and the new value of z_1 is:

$$z_1 + u = z_1 - z_1 + \left(\frac{r}{s}\right)(z_2 - z_1) = \left(\frac{r}{s}\right)(z_2 - z_1).$$

Then, $|z_1 + u| = |\frac{r}{s}(z_2 - z_1)| = \frac{r}{s}|z_2 - z_1| = \frac{r}{s} \cdot s = r$. In contrast, applying the same control to (x_2, m_2) we get the new value of z_2:

$$z_2 + u = z_2 - z_1 + \frac{r}{s}(z_2 - z_1) = \left(1 + \frac{r}{s}\right)(z_2 - z_1).$$

It follows that $|z_2 + u| = |(1 + \frac{r}{s})(z_2 - z_1)| = \left(1 + \frac{r}{s}\right)|z_2 - z_1| = \left(1 + \frac{r}{s}\right)s = s{+}r > r$. Thus the control input u produces $h(f(x_1, m_1), u) = 1 \neq h(f(x_2, m_2), u) = 0$. □

3.4 Continuous-time setting

In Section 5 below, as in Section 2 above, we find it more convenient to work with a continuous-time version of the problem. This does not introduce significant conceptual differences. Instead of the system (2), the dynamics will be described by

$$\dot{x} = f(x, u) \qquad y = h(x)$$

with the meaning of the variables unchanged. Indistinguishable states are still defined to be states that produce identical outputs for any control input. The 1DSLAM example will have dynamics $f(x, u) = u$, the output map remains the same, and the indistinguishable sets for it are the same as for its discrete-time version studied above; we will revisit this case in detail in Section 5.2.

4 Computing Indistinguishability

In this section, we study different approaches for computing indistinguishability relations. Since indistinguishability is a special types of bisimulation, and computing bisimulations is generally hard for nonlinear and hybrid models [10, 11, 26], we can expect that computing indistinguishability relations will also be hard in general. First we show that, under some relatively mild assumptions on f, h, and $\mathcal{U}$, we can decide automatically whether a *candidate* indistinguishability relation $\mathcal{I}$ is indeed so. Next, we present an iterative algorithm that converges to the indistingushable relation, assuming that the computation is performed with the correct hierarchical partitioning of $\mathcal{X} \times \mathcal{X}$. Throughout, we show how these results can be applied to gain insights about the limits of any solution to various versions of the SLAM problem.

4.1 Checking Indistinguishability is Decidable

If f is polynomial, h is a combination of polynomials, absolute values, and $\mathcal{U}$ is a subset of $\mathbb{R}^n$ or or a finite set, then the problem of checking whether a candidate relation $\mathcal{I}$ is indeed an indistinguishablity relation is decidable. This essentially follows from Tarski-Seidenberg Theorem [27].

Proposition 4.1. *If f, h, and the candidate relation $\mathcal{I}$ are described by polynomials and absolute values, indistinguishability of $\mathcal{I}$ is decidable.*

Proof. We can write the existence of distinguishable states in $\mathcal{I}$ as a satisfiability problem: $\exists x_1, x_2 \in \mathcal{X}, u \in \mathcal{U}$

$$(x_1, x_2) \in \mathcal{I} \wedge (h(x_1) \neq h(x_2) \vee (f(x_1, u), f(x_2, u)) \notin \mathcal{I}).$$

Since all the objects involved are expressed as polynomials and modulus functions, the problem is decidable using Tarski-Seidenberg's decision procedure [27]. □

For the nDSLAM example with $n = 1, 2, 3$, using the z3 SMT solver [7] we checked that the corresponding relations defined by Equation (5) are indeed an indistinguishability relations, and a modified relations defined by $x_1 - m_1 = x_2 - m_2 + \mathbf{1}$ are *not*.

4.2 Semi-decision procedure for Indistinguishable Sets

There are two criteria for a pair of states to be related by an indistinguishability relation $\mathcal{I}$: first, $\mathcal{I}$ is closed under control actions, and that the states produce the same observable outputs. One way for computing $\mathcal{I}$ is to iteratively compute its complement $\mathcal{D}$—the distinguishability relation. Informally, we start with $\mathcal{D}_0$ relating state-pairs that are readily distinguishable without control, and then add state pairs to $\mathcal{D}_1$ that become distinguishable with one control input, and so on. In the limit, the union of $\mathcal{D}_0 \cup \mathcal{D}_1 \cup \ldots$ of distinguishable state pairs will give the complement of $\mathcal{I}$. This is the basic idea underlying Algorithm 1. Our implementation of the algorithm is based on the following assumptions.

Assumption 1. (i) The state space $\mathcal{X}$ is a bounded hyperrectangle, (ii) the range of the measurement function $h : \mathcal{X} \to \mathcal{Y}$ is finite, (iii) the input space $\mathcal{U}$ is finite, and (iv) the representation of compact subsets $X, X' \subset \mathcal{X}$ is such that finite unions $X \cup X'$, intersections $X \cap X'$, and preimages under control $f^{-1}(X, u)$ are exactly computable.

The implementation relies on some representation choices which are hidden from the pseudocode for the sake of keeping it uncluttered. We will first discuss the representation issues, before proceeding with the correctness argument.

Representation of relations. For infinite state systems, implementing Algorithm 1 involves choosing representations the relations $\mathcal{I}_k, \mathcal{D}_k, \ldots$, so that Assumption 1 can be satisfied. For a finite-valued measurement function $h : \mathcal{X} \to \mathcal{Y}$, $\mathcal{D}_0$ can be represented as the Cartesian product of the preimages (also called the level sets or fibers) h. That is,

$$\mathcal{D}_0 = \bigcup_{y, y' \in \mathcal{Y}, y \neq y'} h^{-1}(y) \times h^{-1}(y').$$

For the 1DSLAM example, this can be represented by a disjunction of linear inequalities:

$$\mathcal{D}_0 : (|x_1 - m_1| \leq r \wedge |x_2 - m_2| > r) \vee \\ (|x_1 - m_1| > r \wedge |x_2 - m_2| \leq r). \tag{6}$$

Recall that for the example the state space $\mathcal{X}$ is $\mathcal{X} \times \mathcal{M}$. In each subsequent iteration, finite number of Cartesian products of subsets of $\mathcal{X}$ or inequalities involving $\mathcal{X}$ are added to $\mathcal{D}_k$, and therefore, $\mathcal{D}_k$ is always represented by a finite collection of Cartesian products of sets (or inequalities) in $\mathcal{X}$. As for $\mathcal{I}_0 = \mathcal{X} \times \mathcal{X} \setminus \mathcal{D}_0$, we observe that it is an equivalence relation and it can be represented by its partitions. Recall, a *refinement* of R is another equivalence relation R', such that any set B in the partition of R' is also in the partition of R. There are several possible ways of refining $\mathcal{D}_0$ that can satisfy Assumption 1. For concreteness, if $\mathcal{X}$ is a hyperrectangle in $\mathbb{R}^n$, we can consider an uniform grid of $2n$-dimensional hyberrectangles of side r.

The **while** loop continues as long as new distinguishable sets are discovered, i.e., $\mathcal{I}_{k+1}$ is a strict subset of $\mathcal{I}_k$, or the maximum iteration threshold is reached. The **for** loops pick a particular partition $X_1 \times X_2$ from $\mathcal{D}_k$ and a particular control input u from the finite set $\mathcal{U}$, and then computes preimages of X_1 and X_2 under control u and adds them to $\mathcal{D}_{k+1}$. All these newly distinguishable state pairs (if any) are collected in $\mathcal{D}_{k+1}$ Finally, $\mathcal{I}_{k+1}$ is the set that remain after removing $\mathcal{D}_{k+1}$ from $\mathcal{I}_k$. Correctness of Algorithm 1 is stated below.

Proposition 4.2. *For any $k \in \{0, k_{max}\}$, $\mathcal{D}_k$ is the set of k-distinguishable state pairs.*

Proof. The proof is by simple induction on k. For $k = 0$, by construction in Line 2, $\mathcal{D}_0$ is the set of state pairs distinguishable without control inputs. Consider any state pair $(x_1, x_2) \in \mathcal{D}_{k+1}$. From Algorithm 1, it follows that there must be a partition $(X_1, X_2) \in \mathcal{D}_k$ and a control input $u \in \mathcal{U}$ such that $x_1 \in f^{-1}(X_1, u), x_2 \in f^{-1}(X_2, u)$. That is, $(f(x_1, u), f(x_2, u)) \in \mathcal{D}_k$. By the inductive hypothesis, for each $(x_1', x_2') \in \mathcal{D}_k$, there exists $\overline{u} \in \mathcal{U}^k$ such that $h(x_1^{\overline{u}})[k] \neq h(x_2^{\overline{u}})[k]$. Thus, for (x_1, x_2) the control sequence $\overline{u}' = u \cdot \overline{u}$ of length $k+1$ gives $h(x_1^{\overline{u}'})[k+1] \neq h(x_2^{\overline{u}'})[k+1]$ making them $(k+1)$-distinguishable.

Now consider ay $(k+1)$-distinguishable state pair (x_1, x_2). By definition, there must exist (x_1', x_2') that are k-distinguishable, and a control input $u \in \mathcal{U}$ such that $x_1' = f(x_1, u)$ and $x_2' = f(x_2, u)$. By induction hypothesis, $(x_1', x_2') \in \mathcal{D}_k$, that is, they are contained in some partition $X_1 \times X_2$ of $\mathcal{D}_k$, and therefore, by Assumption 1, $(x_1, x_2) \in (f^{-1}(X_1, u), f^{-1}(X_2, u)) \subseteq \mathcal{D}_{k+1}$. □

Algorithm 1: Indistinguishability by iteratively removing distinguishable pairs

Input: Max iterations: $k_{\max}$
Output: Final indistinguishable relation $\mathcal{I}_k$

```
1  Initialize k = 0;
2  D_k = {(x_1, x_2) | h(x_1) ≠ h(x_2)};
3  I_k = X × X \ D_k;
4  I_{k+1} = ∅;
5  while k < k_max and I_{k+1} ⊊ I_k do
6      D_{k+1} = ∅;
7      foreach (X_1, X_2) ∈ D_k do
8          foreach u ∈ U do
9              P = (f^{-1}(X_1, u), f^{-1}(X_2, u)) ∩ I_k;
10             D_{k+1} = D_{k+1} ∪ P;
11     I_{k+1} = I_k \ D_{k+1}; k ← k + 1;
12 return I_k, D_k
```

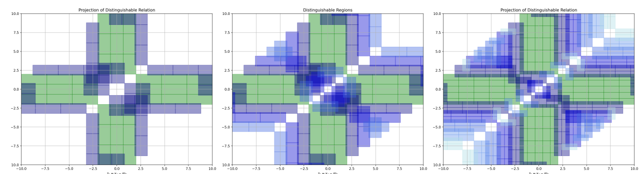

Figure 2: Iterative computation of distinguishable relation $\mathcal{D}_k$ for 1DSLAM projected on the $(x_1 - m_1)$ vs. $(x_2 - m_2)$ plane. $\mathcal{D}_0$ (green) and each successive iteration of $\mathcal{D}_k$ is shown in shades of blue for $\mathcal{U} = \{1, -1\}$. The indistinguishable relation, the diagonal $x_1 - m_1 = x_2 - m_2$ is recovered as complement of union of the $\mathcal{D}_k$'s.

5 Kalman observability decomposition viewpoint

Here we want to interpret and further study the above indistinguishability problem, including the special case of the 1DSLAM example of Section 3, using classical techniques from linear system theory. We can rewrite the example by including the (fixed) landmark location m as a second state, which in continuous time gives the dynamics

$$\dot{x} = u, \quad \dot{m} = 0 \\ y = q(x - m) \tag{7}$$

where

$$q(e) = \begin{cases} 1 & \text{if } |e| \le r \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

for some $r \ge 0$. Alongside this system, we find it helpful to consider the *linear* system obtained by replacing the binary output map q with the identity map:

$$\begin{aligned} \dot{x} &= u, \quad \dot{m} = 0 \\ y &= x - m \end{aligned} \tag{9}$$

Indistinguishable states for this system can be studied with the help of the unobservable subspace and Kalman's well-known observability decomposition (see, e.g., [12, Section 16.1]), which we review below.

For simplicity, in this section we treat the output as being measured continuously in time, without sampling. However, the presence of sampling does not significantly affect the results, as we explain in Remark 2 below.

5.1 Unobservable subspace and Kalman observability decomposition

Consider a general linear system

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx \end{aligned} \tag{10}$$

with $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, and $y \in \mathbb{R}^p$. Without loss of generality, assume that $\operatorname{rank} C = p$, i.e., the map $C : \mathbb{R}^n \to \mathbb{R}^p$ is onto (this can be always be achieved by removing some redundant output components). Define the *observability matrix*

$$\mathcal{O}(A, C) := \begin{pmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{pmatrix}$$

Let $V \subset \mathbb{R}^n$ be the subspace $V := \operatorname{Ker} \mathcal{O}(A, C)$, called the *unobservable subspace*. Let d denote the dimension of V. This subspace has the following properties:

(i) V is contained in the kernel of C;

(ii) V is A-invariant;

(iii) V is the largest (with respect to inclusion) subspace of $\mathbb{R}^n$ with properties (i) and (ii).

Property (i) is obvious from the definitions. Property (ii) is also straightforward to check using Cayley-Hamilton theorem. As for property (iii), an arbitrary subspace V' satisfying (i) and (ii) is seen recursively to be contained in the kernel of $C, CA, \dots,$ hence $V' \subseteq V$.

The subspace V can be generated by the following iterative algorithm: with $c_1, \dots, c_p$ denoting the rows of C, let

$$\begin{aligned} W_0 &:= \operatorname{span}\{c_1, \dots, c_p\} \\ W_k &:= W_{k-1} + \operatorname{span}\{cA : c \in W_{k-1}\}, \qquad k = 1, 2, \dots \end{aligned} \tag{11}$$

This sequence of subspaces stabilizes after at most n steps: $W_n = W_{n-1}$, and we have $W_{n-1}^{\perp} = V$ (the superscript $\perp$ denotes the orthogonal complement).

Claim 1: Given two initial conditions x_0 and x_0' with $x_0 - x_0' \in V$ and an arbitrary control $u(t)$, $t \ge 0$, the two corresponding trajectories of (10) generated by this control produce identical outputs $y(t)$, $t \ge 0$.

As in Section 3, we call such initial conditions *indistinguishable*.

To prove Claim 1, by linearity it is enough to show that the initial condition $\bar{x}_0 := x_0 - x_0'$ under the identically zero input ($u \equiv 0$) produces the identically zero output ($y \equiv 0$). Since $\bar{x}_0 \in V$, by property (ii) of V the corresponding solution satisfies $x(t) = e^{At}\bar{x}_0 \in V$ for all $t \ge 0$, and then by property (i) of V we have $y(t) = Cx(t) = 0$ for all $t \ge 0$, as claimed.

Moreover, if we consider the system obtained from (10) by composing its output with our binary function q from (8):

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= q(Cx) \end{aligned} \tag{12}$$

then it is clear that Claim 1 still holds for this modified system.

Claim 2: Given two initial conditions x_0 and x_0' with $x_0 - x_0' \notin V$ and an arbitrary time $\overline{T} \ge 0$, there exists a control $u(t)$, $t \in [0, \overline{T}]$ such that the two corresponding trajectories of (10) generated by this control produce outputs that are different (at some $t \in [0, \overline{T}]$).

Of course, this control may not be implementable based on the y-measurements. Constructing controls that are actually implementable requires more work, as we know from the 1DSLAM example.

To see why Claim 2 is true, we can revisit the argument that we used to support Claim 1. It is actually enough to consider the zero control $u \equiv 0$. In order for $Ce^{At}\bar{x}_0$ to remain identically zero, $\bar{x}_0$ must belong to an A-invariant subspace contained in the kernel of C. By property (iii) of V, we must then have $\bar{x}_0 \in V$, a contradiction.

A useful way to visualize this situation is to change coordinates in $\mathbb{R}^n$ so that

$$V = \left\{ \begin{pmatrix} x_1 \\ 0 \end{pmatrix} : x_1 \in \mathbb{R}^d \right\} \tag{13}$$

(This simply means picking the first d basis vectors from V and the remaining $n - d$ basis vectors from outside V.) Properties (i) and (ii) of V imply that in the new coordinates, the system takes the form

$$\begin{aligned} \dot{x}_1 &= A_{11}x_1 + A_{12}x_2 + B_1u \\ \dot{x}_2 &= \phantom{A_{11}x_1 +} A_{22}x_2 + B_2u \\ y &= \phantom{A_{11}x_1 + A} C_2x_2 \end{aligned} \tag{14}$$

Moreover, a direct computation shows that we have $\mathcal{O}(A, C) = [0\ \mathcal{O}(A_{22}, C_2)]$ and since this has rank $n - d$ which is the dimension of x_2, the pair (A_{22}, C_2) is observable.

The form (14) is the *Kalman observability decomposition*, and the state components x_1 and x_2 are called the *unobservable* and *observable modes*, respectively. This terminology is clarified by the following observation. Under the zero control $u \equiv 0$, we have

$$\begin{pmatrix} y(0) \\ \dot{y}(0) \\ \vdots \\ y^{(n-1)}(0) \end{pmatrix} = \mathcal{O}(A_{22}, C_2)x_2(0)$$

Since the matrix $\mathcal{O}(A_{22}, C_2)$ has full rank, we see that initial conditions with different x_2-components lead to outputs that are different, even on an arbitrarily small time interval (this is Claim 2). We can in fact uniquely recover $x_2(0)$ from the knowledge of the output function. On the other hand, initial conditions differing only in their x_1-components produce identical outputs (this is Claim 1).

For the modified system (12), which in the above coordinates takes the form

$$\begin{aligned} \dot{x}_1 &= A_{11}x_1 + A_{12}x_2 + B_1u \\ \dot{x}_2 &= \qquad\qquad A_{22}x_2 + B_2u \\ y &= \qquad\quad q(C_2x_2) \end{aligned} \tag{15}$$

it is not clear whether Claim 2 still holds. The previous proof, relying on $u \equiv 0$, will not work, as we will also see below for the 1DSLAM example. So, unlike in the linear case, the roles of the output and the input cannot be decoupled. It thus makes sense to add the following assumption:

Assumption 2. The pair (A_{22}, B_2) is controllable, i.e., we have $\text{rank}(B_2, A_{22}B_2, \cdots, A_{22}^{n-d-1}B_2) = n - d$.

The controllability assumption guarantees that x_2 can always be steered to an arbitrary desired value by some control. A coordinate-free formulation of this assumption is that the controllable subspace of the pair (A, B) and the subspace V together span $\mathbb{R}^n$; in other words, all observable modes are controllable.

Proposition 5.1. *Under the controllability assumption, the statement of Claim 2 is true for the system* (15).

Proof. We let $x(\cdot)$ and $x'(\cdot)$ denote the solutions from the initial conditions x_0 and x_0', respectively. We write $x_1(t), x_2(t)$ for the components of $x(t)$ relative to the coordinates in (14), and similarly we write $x_1'(t), x_2'(t)$ for the corresponding components of $x'(t)$. Apply the solution formula for linear systems to write

$$x_2(t) = e^{A_{22}t}x_2(0) + \int_0^t e^{A_{22}(t-s)}B_2u(s)ds$$

$$x_2'(t) = e^{A_{22}t}x_2'(0) + \int_0^t e^{A_{22}(t-s)}B_2u(s)ds$$

This implies that $C_2x_2'(t) - C_2x_2(t) = C_2e^{A_{22}t}(x_2'(0) - x_2(0))$. By observability, the map $t \mapsto C_2e^{A_{22}t}(x_2'(0) - x_2(0))$ is not identically zero on any time interval $[0, \overline{T}]$ (see the above proof of Claim 2 for the linear case), hence there is some $t \in [0, \overline{T}]$ such that the vector

$$v := C_2e^{A_{22}t}(x_2'(0) - x_2(0))$$

is nonzero. By the controllability assumption, x_2 can be steered to an arbitrary value at this time t, and we also assumed that $p = \text{rank}\, C = \text{rank}(0\; C_2)$ hence the map $C_2 : \mathbb{R}^{n-d} \to \mathbb{R}^p$ is onto. This guarantees the existence of a control that gives $C_2x_2(t) = (r - \varepsilon)v/|v|$ for some $\varepsilon \in (0, |v|)$. We have $|C_2x_2(t)| = r - \varepsilon < r$ hence $q(C_2x_2(t)) = 0$. On the other hand, $C_2x_2'(t) = C_2x_2(t) + C_2(x_2'(t) - x_2(t)) = C_2x_2(t) + C_2e^{A_{22}t}(x_2'(0) - x_2(0)) = (r - \varepsilon)v/|v| + v$ and its norm is $|C_2x_2'(t)| = r - \varepsilon + |v| > r$, hence $q(C_2x_2'(t)) = 1$. Since the two output values are different, the claim is established. □

Remark 1. An explicit expression for the distinguishing controller in Proposition 5.1 can be obtained in three steps. First, a time $t \in [0, \overline{T}]$ can in practice be picked arbitrarily, because all but a finite number of values for t will give $v \neq 0$. Second, we need to find an $x_2(t)$ satisfying $C_2x_2(t) = (r - \varepsilon)v/|v|$; one such choice is $x_2(t) = C_2^T(C_2C_2^T)^{-1}(r - \varepsilon)v/|v|$, where the inverse exists by our assumption that the rows of C_2 are linearly independent. Finally, letting $W(0, t) := \int_0^t e^{A_{22}s}B_2B_2^Te^{-A_{22}^Ts}ds$ denote the controllability Gramian of the pair (A_{22}, B_2), which is invertible by our controllability assumption, a desired control achieving the transfer from $x_2(0)$ to $x_2(t)$ is

$$u(s) := B_2^Te^{-A_{22}^Ts}W^{-1}(0, t)\left[e^{-A_{22}t}x_2(t) - x_2(0)\right], \quad s \in [0, t]$$

(this is verified by direct computation; see [3, Section 13]). Of course, this controller requires precise knowledge of the initial values $x_2(0)$ and $x_2'(0)$ in order to compute the vector v, and as such it is not implementable with output data only.

Remark 2. We can also consider the situation where, as before, the output is measured at sampling times kT, $k = 0, 1, \ldots$ rather than continuously. Clearly, Claim 1 remains valid in the presence of sampling. Claim 2 for the linear system (10) remains valid for almost all values of the sampling period T; this follows from the well-known result (see, e.g., [25, Section 6.2]) that observability of linear systems is preserved under sampling as long as the sampling period satisfies a suitable "non-resonance" condition. As for Proposition 5.1, which establishes Claim 2 for the system with binary output measurements, it also remains valid for almost all values of T, for the same reason just given and also because, as explained in Remark 1, for almost all T we can transfer x_2 to the desired value exactly at T (or another integer multiple of T).

5.2 Back to the 1DSLAM example

We can now return to the system (7) describing the 1DSLAM example, and its linear counterpart (9). The linear system data is: $A = 0$, $B = \binom{1}{0}$, $C = (1\; -1)$. The subspace algorithm (11) gives $W_0 = \text{span}\{(1\; -1)\} = W_1$, hence $V = W_0^\perp = \text{span}\{\binom{1}{1}\}$. Claim 1 now says that pairs (x, m), (x', m') are indistinguishable when $\binom{x}{m} - \binom{x'}{m'} \in \text{span}\{\binom{1}{1}\}$, which is equivalent to $x - m = x' - m'$. In other words, any two points on a diagonal in the (x, m)-space corresponding to a fixed relative distance between the agent and the landmark are indistinguishable. For the discrete-time version of the example, we already established this fact directly in Proposition 3.7. The fact that Claim 2 is valid for this system was already shown (for the discrete-time case) in Proposition 3.8.

We can apply a simple coordinate change[2] $(x, m) \mapsto (x, e)$ with $e := x - m$ to arrive at the dynamics

$$\begin{aligned} \dot{x} &= u, \quad \dot{e} = u \\ y &= q(e) \end{aligned} \tag{16}$$

A special case of (15), this is the Kalman observability decomposition for this example. In the new coordinates, Claim 1 says that two points (x, e) and (x', e') are indistinguishable when $e = e'$. With regards to Claim 2, here we see again the difference with the linear case in that the choice of control is nontrivial. The identically zero control will not work because e stays constant under such control, and so any two initial states with the same output will lead to trajectories with identical outputs. On the other hand, the scalar system $\dot{e} = u$ is controllable, and this property (or rather its discrete-time counterpart) was implicitly used in proving Proposition 3.8. (The

[2] We note that the condition (13) does not determine a coordinate change uniquely.

construction there is actually very similar to the one in the proof of Proposition 5.1.)

Earlier (see the proof of Proposition 3.8), working with the discrete-time variant of the 1DSLAM example, we derived an *implementable* controller that steers e to 0 (or to any other desired value), thus narrowing down an indistinguishable set to which the state belongs. It remains to be seen whether the Kalman observability decomposition, coupled with the controllability assumption, can help us systematically design such controllers. We also need to investigate whether the controllability assumption can be relaxed.

5.3 Nonlinear systems

The above approach can also be extended to nonlinear dynamics, using tools from geometric nonlinear control theory (see, e.g., [13]). Consider a system in the form

$$\dot{x} = f(x) + g(x)u, \quad y = h(x) \tag{17}$$

We take the output map $h(\cdot)$ to be smooth; our goal here is to establish a counterpart of Claim 1 from Section 5.1, and we know that such a claim will automatically apply to a quantized version of the output. To simplify the discussion, we also assume that both the input u and the output y are scalar (although this is not necessary).

According to the results in [13, Section 1.9], we should look for the largest *distribution* (i.e., a smooth assignment of a subspace to each state x) that is: (i) involutive (closed under Lie bracketing); (ii) contained in the distribution $(\text{span}\{dh\})^{\perp}$ (with dh being the row vector of partial derivatives of the output map h with respect to the state variables, and $\perp$ denoting the orthogonal complement as before); and (iii) invariant under the vector fields f and g. We will not formally define all the concepts just mentioned; instead, we refer the reader to [13] for the definitions, but will work out an example below to illustrate the ideas. For constructing such a distribution, instead of working with condition (ii) above it is more practical to look for the smallest codistribution[3] containing $\text{span}\{dh\}$ and invariant under f and g, and then take the orthogonal complement to arrive at the desired distribution. An algorithmic procedure for doing this is given in [13, Section 1.9], and it is a direct nonlinear extension of the iterative algorithm (11).

Without going into a detailed description of the general procedure, we illustrate how it works on a specific example. Consider the system

$$\dot{x} = v, \quad \dot{v} = u - kv^2, \quad \dot{m} = 0$$
$$y = x - m$$

This system can be considered as a special case of the linear system (9) that we treated earlier, augmented with an additional velocity state v and incorporating a drag force term, with k being a drag coefficient. We can rewrite this system in the form (17) as

$$\begin{pmatrix} \dot{x} \\ \dot{v} \\ \dot{m} \end{pmatrix} = \begin{pmatrix} v \\ -kv^2 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} u =: f(x, v, m) + gu$$

$$y = \begin{pmatrix} 1 & 0 & -1 \end{pmatrix} \begin{pmatrix} x \\ v \\ m \end{pmatrix}$$

[3]For our purposes, we can think of a codistribution as an object analogous to a distribution except that it is spanned by row vectors instead of column vectors.

Let c denote the row vector $\begin{pmatrix} 1 & 0 & -1 \end{pmatrix}$. The algorithm starts with the codistribution $\text{span}\{dh\} = \text{span}\{c\}$. The fact that this is a constant codistribution simplifies the subsequent calculations. Next, the following quantities must be computed:

$$L_f c = L_f(dh) = d(L_f h) = d\Big(\frac{\partial h}{\partial x} \cdot f\Big) = dv = \begin{pmatrix} 0 & 1 & 0 \end{pmatrix},$$
$$L_g c = 0 \ \text{(since } g \text{ is constant)},$$
$$L_f(L_f c) = L_f(dv) = d(L_f v) = d(-kv^2) = \begin{pmatrix} 0 & -2kv & 0 \end{pmatrix}$$

We see that the row vector obtained in the last computation is linearly dependent on $L_f c$, hence the procedure stops and yields the (constant) codistribution

$$\text{span}\{\begin{pmatrix} 1 & 0 & -1 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 \end{pmatrix}\}$$

spanned by the two linearly independent row vectors collected during the procedure. Its orthogonal complement

$$V = \text{span}\{ \begin{pmatrix} 1 & 0 & 1 \end{pmatrix}^T \} \tag{18}$$

is the desired distribution; in fact, in this example it is a constant distribution, i.e., just a linear subspace. It can be directly checked that V satisfies the properties (i)–(iii) listed earlier; in particular, it is orthogonal to c and commutes with (hence is invariant under) both f and g. Hence, the conclusion for this example is:

Initial conditions whose difference lies in the subspace V given by (18) *are indistinguishable.*

Intuitively this conclusion makes sense and is consistent with our findings for the simpler version (9) of this example.

A major open question in this context is whether a version of the Claim 2 can be obtained, especially when the output is composed with the binary function (8). In other words, are pairs of points whose difference is not in V actually distinguishable (by some control, and under suitable controllability assumptions) based on such output measurements?

It is important to note that the method described in this section relies on differentiability of the vector fields involved (f, g, h) and as such is not readily applicable to hybrid systems. In this respect, the computational approach of Section 4 (or some combination of the two) may hold more promise.

6 Conclusions

We investigated a version of the SLAM problem in which an agent has to localize itself relative to an unknown landmark using finite-valued measurements. To understand the fundamental limits imposed by such coarse measurements, we focused on characterizing *indistinguishable states*. We demonstrated that indistinguishability relations can be verified automatically and that they can be approximated iteratively. We presented an alternative analytical approach based on the observability theory of linear control systems, which iteratively computes the indistinguishable subspace in finitely many steps. A differential-geometric extension for nonlinear systems was also developed.

One direction for future work is to relate indistinguishability more precisely to fundamental limits for the SLAM problem. The size of the indistinguishable set gives us a lower bound on the size of the *initial* joint uncertainty. It is possible, however, that this uncertainty reduces over time (e.g., if the system dynamics are contracting) even without the needing further measurements. In the examples we considered, this did not happen as the system dynamics were volume-preserving. General assumptions under which the size of uncertainty at any future time can be characterized in terms of indistinguishable sets remain to be formalized.

Furthermore, knowledge of indistinguishable sets translates to the limits on achievable control tasks. If one of the two indistinguishable initial states is known to lead to unsafe trajectories, then assuring safe control is hopeless. More generally, the size and shape of indistinguishable sets relates to the size and shape of achievable target sets, although this relationship is nontrivial and complete characterization should be investigated further.

References

[1] J. Andrade-Cetto and A. Sanfeliu. The effects of partial observability in SLAM. In *Proc.* IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), pages 397–402, 2004.

[2] P. Bernard. *Observer Design for Nonlinear Systems*, volume 479 of *Lecture Notes on Control and Information Sciences*. Springer, 2019.

[3] R. W. Brockett. *Finite Dimensional Linear Systems*. Wiley, New York, 1970.

[4] M. Buchanan. Leibniz's principle of indistinguishability. *Nature Physics*, 15(11):1104–1104, 2019.

[5] Y. Cheng, M. Maimone, and L. Matthies. Visual odometry on the mars exploration rovers. In *2005 IEEE International Conference on Systems, Man and Cybernetics*, volume 1, pages 903–910. IEEE, 2005.

[6] S. Clarke and G. W. F. von Leibniz. *The Leibniz-Clarke Correspondence: Together Wiith Extracts from Newton's Principia and Opticks*. Manchester University Press, 1956.

[7] L. De Moura and N. Bjørner. Z3: an efficient SMT solver. In *Proceedings of the Theory and practice of software, 14th international conference on Tools and algorithms for the construction and analysis of systems*, TACAS'08/ETAPS'08, pages 337–340, Berlin, Heidelberg, 2008. Springer-Verlag.

[8] M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba. A solution to the simultaneous localization and map building (slam) problem. *IEEE Transactions on robotics and automation*, 17(3):229–241, 2001.

[9] J. Engel, V. Koltun, and D. Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2017.

[10] A. Girard and G. J. Pappas. Approximate bisimulations for nonlinear dynamical systems. In *Proc. of IEEE Conference on Decision and Control*, pages 684–689. IEEE, 2005.

[11] T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What's decidable about hybrid automata? In *ACM Symposium on Theory of Computing*, pages 373–382, 1995.

[12] J. P. Hespanha. *Linear Systems Theory*. Princeton University Press, 2009.

[13] A. Isidori. *Nonlinear Control Systems*. Springer, Berlin, 3rd edition, 1995.

[14] A. Katok and B. Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, 1995.

[15] M. Khaled, K. Zhang, and M. Zamani. A framework for output-feedback symbolic control. *IEEE Transactions on Automatic Control*, 68(9):5600–5607, 2022.

[16] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.

[17] K. W. Lee, W. S. Wijesoma, and J. I. Guzman. On the observability and observability analysis of SLAM. In *Proc.* IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), pages 3569–3574, 2006.

[18] D. Liberzon and S. Mitra. Entropy and minimal bit rates for state estimation and model detection. *IEEE Trans. Automat. Control*, 63:3330–3344, 2018.

[19] N. Lynch. A hundred impossibility proofs for distributed computing. In *Proceedings of the eighth annual ACM Symposium on Principles of distributed computing*, pages 1–28, 1989.

[20] R. Majumdar, N. Ozay, and A.-K. Schmuck. On abstraction-based controller design with output feedback. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, pages 1–11, 2020.

[21] R. Milner. *A calculus of communicating systems*. Springer, 1980.

[22] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans. Feedback control under data rate constraints: An overview. *Proc. IEEE*, 95:108–137, 2007.

[23] D. Park. Concurrency and automata on infinite sequences. In *Theoretical Computer Science: 5th GI-Conference Karlsruhe, March 23–25, 1981*, pages 167–183. Springer, 2005.

[24] H. Sibai and S. Mitra. State estimation of continuous-time dynamical systems with uncertain inputs with bounded variation: Entropy, bit rates, and relation with switched systems. *IEEE Trans. on Automatic Control (TAC)*, 68(12):7041–7056, 2023.

[25] E. D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer, New York, 2nd edition, 1998.

[26] P. Tabuada and G. J. Pappas. Finite bisimulations of controllable linear systems. In *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, volume 1, pages 634–639. IEEE, 2003.

[27] A. Tarski. A decision method for elementary algebra and geometry. In *Quantifier elimination and cylindrical algebraic decomposition*, pages 24–84. Springer, 1998.

[28] S. Thrun, W. Burgard, and D. Fox. *Probabilistic robotics*. MIT Press, 2005.

[29] B. Tovar, A. Yershova, J. M. O'Kane, and S. M. LaValle. Information spaces for mobile robots. In *Proceedings of the Fifth International Workshop on Robot Motion and Control, 2005. RoMoCo'05.*, pages 11–20. IEEE, 2005.

[30] V. K. Weinstein and S. M. LaValle. Equivalent environments and covering spaces for robots. In *Topology and AI*, pages 31–64. ESIAM, 2024.

[31] R. W. Wolcott and R. M. Eustice. Visual localization within lidar maps for automated urban driving. In *2014* IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 176–183. IEEE, 2014.

A Appendix: Irrational shifts and windings

For arbitrary $\alpha, L \in \mathbb{R}$, consider the map from the interval $[0, L]$ to itself defined as

$$x \mapsto x + \alpha \pmod{L} \tag{19}$$

We'll say that α and L are *rationally independent* if

$$\alpha m \neq Ln \quad \forall\, m, n \in \mathbb{Z} \tag{20}$$

A consequence of (20) for the map (19) is that an orbit of the map never hits the same point twice. The following fact is well known. (See, e.g., [14, Section 1.4] for a slightly different proof.)

Lemma 1. *Under the condition* (20)*, every orbit of the map* (19) *is everywhere dense in the interval* $[0, L]$.

Proof. Starting from an arbitrary initial point, apply the map sufficiently many times to wrap around the interval once. This generates a lattice with the distance between neighboring points not exceeding α. By (20), the points on the orbit never repeat. So, during the second wrap-around, there will be two neighboring points on the orbit with distance not exceeding (actually strictly less than) $\alpha/2$. We see that a certain number of iterations of the map (19) generates a shift by less than $\alpha/2 \pmod{L}$. Therefore, if we keep applying the map, the resulting orbit will contain a lattice covering the interval with distance between neighboring points less than $\alpha/2$. Repeating this argument, we can show that the orbit contains a lattice of points separated by less than $\alpha/4$, $\alpha/8$, etc. Thus the orbit is everywhere dense. □

We mention the following well-known corollary of Lemma 1. Here we take the torus to be the square $[0, 1] \times [0, 1]$ with the opposite points identified in the usual way. The same argument works with trivial changes if we take the square $[0, L] \times [0, L]$, for an arbitrary L.

Corollary 1. *A line with an irrational slope* α *densely fills the torus.*

Proof. Take an arbitrary point (x, y) on the torus. We need to show that we can get arbitrarily close to it by following the given line with slope α. With no loss of generality, assume the line passes through the point $(0, 0)$. Then the point $(x, x\alpha)$ is on the line. Each winding returns the horizontal position to x and adds α to the vertical position. This corresponds to the shift map $z \mapsto z + \alpha$ (mod 1) on the vertical interval $[0, 1]$. Since α is irrational, α and 1 are rationally independent in the sense of (20). By Lemma 1, the orbit of this map contains points arbitrarily close to y, as needed. □

For a general L, the vertical shift is $L\alpha$ which is rationally independent from L, so the argument still works.