

Learned Risk Metric Maps for Kinodynamic Systems

Ross E. Allen¹, Wei Xiao², and Daniela Rus²

Abstract—We present Learned Risk Metric Maps (LRMM) for real-time estimation of coherent risk metrics of high-dimensional dynamical systems operating in unstructured, partially observed environments. LRMM models are simple to design and train—requiring only procedural generation of obstacle sets, state and control sampling, and supervised training of a function approximator—which makes them broadly applicable to arbitrary system dynamics and obstacle sets. In a parallel autonomy setting, we demonstrate the model’s ability to rapidly infer collision probabilities of a fast-moving car-like robot driving recklessly in an obstructed environment; allowing the LRMM agent to intervene, take control of the vehicle, and avoid collisions. In this time-critical scenario, we show that LRMMs can evaluate risk metrics 20-100x times faster than alternative safety algorithms based on control barrier functions (CBFs) and Hamilton-Jacobi reachability (HJ-reach), leading to 5-15% fewer obstacle collisions by the LRMM agent than CBFs and HJ-reach. This performance improvement comes in spite of the fact that the LRMM model only has access to local/partial observation of obstacles, whereas the CBF and HJ-reach agents are granted privileged/global information. We also show that our model can be equally well trained on a 12-dimensional quadrotor system operating in an obstructed indoor environment. The LRMM codebase is provided at <https://github.com/mit-drl/pyrmm>.

I. INTRODUCTION

The estimation of risk is a topic that spans many research domains: from medical treatments [1], to economics and portfolio optimization [2], [3], to robotics and autonomous systems [4]. Within the field of autonomous systems, we often consider risk in terms of probability that a system causes harm to itself or its environment, and the expected cost of such events. In this work we seek to estimate the probability of a dynamical system arriving at a state of failure given an initial configuration and control policy.

Such risk estimation problems have been a major focus within autonomous automobile research where collision avoidance is of utmost concern [6]. Many of these works emphasise prediction of future states of other vehicles [7], [8], [9]. These works often adopt domain-specific assumptions that make it difficult to extend techniques beyond the

Research was sponsored by the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

¹Massachusetts Institute of Technology, Lincoln Laboratory, Lexington, MA, 02421, USA ross.allen@ll.mit.edu

²Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, 02139, USA weixy@mit.edu; rus@csail.mit.edu

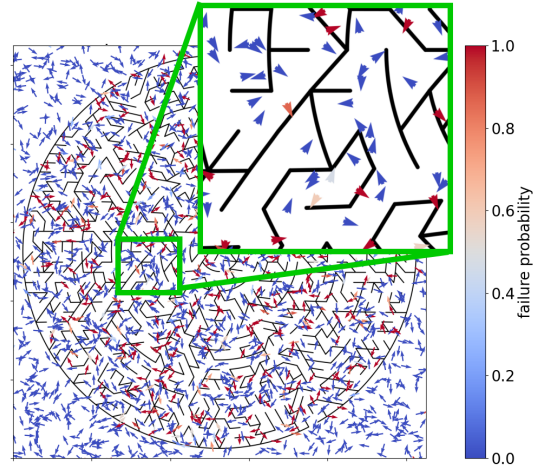


Fig. 1. Visualization of inferred risk metrics for 1000 sampled states of a Dubins vehicle within a procedurally generated maze [5]. The minimum turning radius of the vehicle is ~ 10 pixels. The estimated probability of collision with walls for a random control policy over a finite time horizon is colorized. We see that states with an imminent wall collision are correctly estimated as high risk, whereas unobstructed states are low risk of collision. For clarification, the base of each arrow, not its tip, represent the exact xy -coordinate of the sampled state.

automotive domain—such as strict assumptions on vehicle dynamics[8]; the assumption that all obstacles or other agents can be represented as points [8], balls [9], or ellipsoids [7]; and assumptions that other agents act in predictable [7], [9] and/or self-preserving [8] fashions.

A closely related problem is that of trajectory collision probability assessment [10], [11] and risk-aware motion planning [10], [12], [13]. These works evaluate the probability that a particular trajectory intersects the obstacle space given bounded process and/or measurement noise. This is a related—yet *distinct*—problem from the one considered in this paper: i.e. evaluating the probability that a system will arrive at a future failure state given some initial state and control policy. Furthermore risk-aware motion planning algorithms almost all rely on an *a priori* map of the obstacle space with, at most, bounded uncertainty on obstacle locations [12]. Therefore, even though our work shares concepts and keywords with the field of sampling-based motion planning, it occupies a different problem space than algorithms like rapidly-exploring random tree (RRT) [14].

Hamilton-Jacobi reachability (HJ-reach) analysis [15] and control barrier functions (CBFs) [16]—both of which are discussed more thoroughly in Sec. II and IV—offer rigorous methods for estimating unsafe sets and providing safety-guaranteed controls for autonomous systems. However, the former suffers from limitations on scaling to high-

dimensional systems (i.e the “curse of dimensionality”) and the latter lacks generalizability to arbitrary system dynamics and obstacle sets. The scalability limitations of HJ-reachability and generalizability limitations of CBFs motivate the need for alternative methods for assessing safety of dynamical systems and discovering safe control policies. Drawing from the field of motion planning, sampling-based methods are a well-known technique for overcoming the curse of dimensionality [14]. Sampling-based techniques can be applied to systems with arbitrary dynamics and obstacle spaces since these properties are inherent to the sampling process. The primary drawback of sampling-based techniques to safety-critical applications is that they may lack the strict safety guarantees that alternative methods provide. Lew et al. [17], [18] provide strong justification for the use of sampling-based methods for safety-critical applications by proving asymptotic convergence to a conservative over-approximation of reachable sets using random set theory.

In this paper we provide a simple, sampling-based method for approximating risk metrics for arbitrary dynamical systems in partially-observed environments and train a supervised-learning model—referred to as learned risk metric maps (LRMM)—that can rapidly infer failure probabilities in *a priori* unknown environments. We then show that, in spite of HJ-reach and CBF safety guarantees, the LRMM model can provide a greater level of safety in a time-critical parallel autonomy task due to its rapid estimation of risk.

II. RELATED WORK

In this section we position our work relative to the closely related techniques of Hamilton-Jacobi reachability (HJ-reach) and control barrier functions (CBF).

The theoretical foundations for this paper can be traced back to the concept of inevitable collision states; i.e. the set of states for which a dynamical system cannot avoid future collisions with obstacles regardless of control input [19], [20]. HJ-reachability represents a modern treatment of inevitable collision states that provides a rigorous theoretical foundation for computing the set of unsafe states that may lead to failure (e.g. collision); often referred to as the “backward reachable set” (BRS) [15].

The BRS of the obstacle space, \mathcal{C}_{obs} , is computed by solving for the value function, $V(t, s_\omega)$, of the Hamilton-Jacobi-Isaacs partial differential equation (PDE) for all states over time horizon t [21], [22]. The value function then allows us to describe the BRS of the obstacle space as $\mathcal{V}_{\text{obs}}(t) = \{s : V(t, x) \leq 0\}$. This is the set—often called the *unsafe set*—for which there exists a control action, a that would lead the system to collision with the obstacle space over time horizon t . The HJ value function also provides a means to calculate the optimal control for avoiding obstacles as a gradient of the value function; see Bajcsy et al. for details [23].

HJ-reachability has been used to develop provably-safe autonomous navigation algorithms for arbitrary dynamical systems within partially-observed environments by treating the unknown portion of the environment as an obstacle [23].

While the general applicability of HJ-reachability makes it a powerful tool, it suffers from the “curse of dimensionality” that make it very difficult to apply the technique in real-time on anything save the simplest, low-dimensional, slow-moving systems [15], [24], [23].

Control Barrier Functions (CBF) offer an alternative, complementary approach to HJ-reachability that provide provably-safe control of autonomous systems in the presence of obstacles [16], [25]. In contrast to HJ-reachability, CBFs can provide real-time safe control for high-dimensional control-affine systems. CBF-based controllers—or more precisely, controllers based on CBFs *and* control Lyapunov functions (CLFs)—work by mapping state (safety) constraints onto a set of control constraints by taking Lie derivatives of the constraints along the dynamics, and then encoding them as inequality constraints within a quadratic program (QP) that can be solved to determine stabilizing (i.e. target-tracking, using CLFs) and safe (i.e. obstacle-avoiding, using CBFs) control inputs; see Ames et al. [16] for a detailed discussion.

General methods do not exist for discovering CBFs for a given control system, requiring that they be hand-designed. Furthermore CBFs require an explicit mathematical model of obstacles or keep-out regions to be avoided; something that is often impractical in real-world applications. Choi et al. [24] propose a method for blending HJ-reachability and CBFs; however this technique still suffers from the poor scalability to high-dimensional systems.

The learned risk metric maps (LRMM) described in the following sections attempt to overcome the drawbacks of HJ-reach and CBF methods by providing a constructive (i.e. not hand-designed) model that rapidly estimates risk metrics for high-dimensional dynamical systems within arbitrary obstacle spaces.

III. RISK METRIC MAPS

Our work considers partially observable Markov decision processes [26] defined as the tuple $(\mathcal{C}, \mathcal{A}, \mathcal{O}, T, R)$. \mathcal{C} is the configuration space that defines the possible states of the system, $s \in \mathcal{C}$, as well as the obstacles and/or failure-states of the system, $\mathcal{C}_{\text{obs}} \subset \mathcal{C}$. Let the obstacle-free set of the configuration space be defined as $\mathcal{C}_{\text{free}} = \mathcal{C} \setminus \mathcal{C}_{\text{obs}}$. \mathcal{A} is the action space and an action is given as $a \in \mathcal{A}$. $\mathcal{O}(s)$ is the observation function; an observation is given as $o \sim \mathcal{O}(s)$. The state transition function, $T(s'|s, a)$, represents probability of arriving in state s' when taking action a in state s . The reward is drawn from the reward function $r \sim R(s, a)$. A stochastic policy $\pi(a|o) \in \Pi$ represents the probability, or probability density, of taking action a given observation o .

An observation-action trajectory from timestep α to ω is defined as $\tau_{\alpha, \omega} = (o^{(\alpha)}, a^{(\alpha)}, o^{(\alpha+1)}, a^{(\alpha+1)}, \dots, o^{(\omega-1)}, a^{(\omega-1)}, o^{(\omega)})$. Note that we use parenthetical superscripts to represent specific time steps. Let \mathcal{T} be the set of all possible trajectories τ under dynamics T and policy π . Let $J(\tau_{\alpha, \omega}) : \mathcal{T} \rightarrow \mathbb{R}$ be the *trajectory cost function* that maps a trajectory to a real-valued number (e.g. fuel usage, time, etc.).

With slight abuse of notation on the variables α and ω , let us define the *cost-limited forward reachable set* [27] from state s_α to all states s_ω

$$\Omega(s_\alpha, J_{\text{th}}) = \{s_\omega \in \mathcal{C} \mid \exists \tau_{\alpha,\omega} \in \mathcal{T}, J(\tau_{\alpha,\omega}) \leq J_{\text{th}}\} \quad (1)$$

Define the *failure cost function*

$$Z : \mathcal{C} \rightarrow \mathbb{R} \quad (2)$$

that maps each state $s \in \mathcal{C}$ to a real value that corresponds to the failure-state set \mathcal{C}_{obs} . In this work we define a failure cost function such that

$$\begin{aligned} Z(s_\omega \in \mathcal{C}_{\text{obs}}) &= 1 \\ Z(s_\omega \in \mathcal{C}_{\text{free}}) &= 0 \end{aligned} \quad (3)$$

Let \mathcal{Z} be the set of all failure cost functions Z . Define the *risk metric* as

$$\rho : \mathcal{Z} \rightarrow \mathbb{R} \quad (4)$$

In this work we only consider *coherent risk metrics* such as conditional value at risk (CVaR), worst case, or expected cost [4].

Finally—with slight abuse of notation on ρ —we can define the *risk metric map*

$$\rho(s, \pi; Z, \mathcal{C}_{\text{obs}}, T) : \mathcal{C} \times \Pi \rightarrow [0, 1] \quad (5)$$

which maps the configuration and policy space to the real-value range $[0, 1]$ by assigning to each state-policy pair (s, π) a risk value parameterized by the failure cost function Z , obstacle space \mathcal{C}_{obs} , and vehicle dynamics T . By using the binary failure cost function in Eqn. 3 and expected cost as our risk metric, our risk metric map is identical to the *probability* the system arrives at a failure state over an infinite time horizon when starting from state s and following policy π . For brevity of notation we drop the parameter variables and refer to the risk metric map as $\rho(s, \pi)$ or even $\rho(s)$ when the policy is assumed to be uniform random sampling of bounded controls. Note that we can relate the risk metric map to *inevitable collision obstacles* (ICO) [19] as

$$\text{ICO}(\mathcal{C}_{\text{obs}}) = \{s \in \mathcal{C} \mid \forall \pi, \rho(s, \pi) = 1\} \quad (6)$$

A. Approximate Risk Metrics

Many kinodynamic systems of interest—such as cars, aircraft, maritime vessels, etc.—have continuous state and action spaces making enumeration over all states and control trajectories in Eqn. 1 impossible. Furthermore, obstacle spaces in Eqn. 3 are often implicitly defined and only sensed with partial observability. This means that the explicit derivation of risk metrics at each state of a system is impractical, if not impossible.

We can, however, formulate a finite-horizon approximation of Eqn. 5 at a given state. Algorithm 1 gives a recursive, sampling-based estimate of the risk metric at state s , which is arrived at by trajectory τ , and subsequently following policy π over a time horizon of $t * m$.

The algorithm works by checking if a state is in collision with obstacles (i.e. failure cost function from Eqn. 3) and

Algorithm 1 Approximate Risk Metric

```

procedure APPRXRISKMETRIC( $s, \tau, \pi, t, n, m$ )
   $z \leftarrow \text{CHECKFAILURE}(s, \tau)$ 
  if  $z = 1$  or  $m = 0$  then
    return  $z$ 
   $P \leftarrow \emptyset$ 
   $V, E \leftarrow \text{SAMPLEFORWARDREACHABLESET}(s, \pi, t, n)$ 
  for  $s_\omega, \tau_\omega$  in ZIP( $V, E$ ) do
     $P \leftarrow P \cup \{\text{APPRXRISKMETRIC}(s_\omega, \tau_\omega, \pi, t, n, m - 1)\}$ 
  return COHERENTRISKMETRIC( $P$ )

```

then recursively sampling the stochastic policy π which samples the forward reachable set $\Omega(s, t)$ of state s with a cost function of time t (Eqn. 1). This recursive sampling generates a tree rooted at s with a branching factor of n and depth of m . The recursive base of the algorithm is to simply return the value of the failure cost function at the leaf node, which is either 0 or 1. Otherwise the algorithm returns a coherent risk metric—such as expected cost—over all immediately adjacent sampled states in the tree.

B. Learned Risk Metric Maps

While Alg. 1 provides an approximation to risk metrics in Eqn. 5, this only provides an approximation at a single state and not an approximation over the entire state space. Furthermore, we need privileged access to obstacle information in order to perform collision checking during the sampling process. This means that Alg. 1 is of limited use in real-time systems where obstacles may only be partially sensed (e.g. via cameras or Lidar).

We seek a method for approximating risk metrics over finite time horizons for kinodynamic systems operating in unstructured, partially observed obstacle spaces. To achieve this we use Alg. 1 as a data generator to train a function approximator (e.g. neural network) that takes as input the local observation, $o \sim \mathcal{O}(s)$, and outputs the *inferred* risk metric, $\hat{\rho}(s)$. Training data is generated under policy π_g (e.g. uniform random from bounded control space [17], [18]). Therefore we define the *learned risk metric map* (LRMM) function approximator as

$$\hat{\rho}(o; \pi_g) : \mathcal{O} \rightarrow [0, 1] \quad (7)$$

Algorithm 2 gives a rough overview of the relatively simple procedure for training the LRMM model in Eqn. 7. This consists of generating obstacle sets; sampling states from these obstructed configuration spaces; pulling observations from these states and computing risk metrics at each state; and finally, training a supervised model on these observations and risk metrics.

Similar to sampling-based motion planners (e.g. RRT [14]) and online planning methods (e.g. monte carlo tree search [26]) approximate risk metrics use sampling techniques to explore the configuration space using tree structures. However, LRMM differs in that it does not directly

Algorithm 2 Training Learned Risk Metric Maps

```
 $\mathbf{C}_M \leftarrow \text{GENERATEOBSTACLESETS}(M)$   
 $\mathbf{S}_N \leftarrow \text{SAMPLESTATES}(\mathbf{C}_M, N)$   
 $\mathbf{O}_N \leftarrow \text{OBSERVESTATES}(\mathbf{S}_N)$   
 $\mathbf{P}_N \leftarrow \text{APPRXRISKMETRICS}(\mathbf{S}_N, \tau, \pi, t, n, m)$   
 $\hat{\rho} \leftarrow \text{TRAINMODEL}(\text{inputs}=\mathbf{O}_N, \text{targets}=\mathbf{P}_N)$ 
```

attempt to determine policies or motion plans for navigating the configuration space. Therefore, it is not necessary to form a connected graph with all sampled states and instead allows for many disjoint trees sampled throughout the configuration space making it highly parallelizable.

IV. EXPERIMENTS

In this section we provide simulation experiments demonstrating the utility of learned risk metric maps. We provide a case study on parallel autonomy, comparing effectiveness of LRMMs with control barrier functions (CBFs) and Hamilton-Jacobi reachability (HJ-reach) in a 4-dimensional car-like robot. We then provide training results for high-dimensional systems in unstructured obstacles spaces such as a quadrotor in an obstructed room.

For all experiments the LRMM models are trained with Alg. 2 using a branching factor $n = 32$, a tree depth $m = 2$, time horizon of $t = 2$ seconds per tree depth, and expected value of our failure cost function (Eqns. 3) as our coherent risk metric. Similarly, a held-out test set of data is generated. Models consist of a 64-neuron, single layer, feed-forward network with a sigmoid activation function. Another sigmoid function is applied at the output layer to bound outputs to the range $[0, 1]$ so that outputs represent probabilities. The training process uses a mean squared error (MSE) loss function for the regression problem.

For software implementation, state and control spaces are defined and sampled using Open Motion Planning Library (OMPL) [28]. Network models and training are implemented with PyTorch [29]. In Sec. IV-A CBF quadratic programs are solved using the CVXOPT library [30]. The HJ-reachability PDE is defined and solved with the OptimizedDP library [31]. In Sec. IV-B PyBullet [32] is used for quadrotor system dynamics, 3D collision checking, and rendering. All data generation, model training, and evaluation experiment configurations are managed using hydra-zen [33]. LRMM training and experiment software is provided at <https://github.com/mit-drl/pyrmm>.

A. Parallel Autonomy Case Study

Problem setup. We consider a Dubins-like kinodynamic system—referred to as Dubins4D and illustrated in Fig. 2—which is commonly studied in related works [23], [25] with dynamics given as

$$\dot{x} = v \cos \theta + d_x, \quad \dot{y} = v \sin \theta + d_y, \quad \dot{\theta} = u_1, \quad \dot{v} = u_2 \quad (8)$$

where state, $s := (x, y, \theta, v)$, represents the location in Cartesian frame, heading, and linear speed of the vehicle, respectively; and controls, $\mathbf{u} := (u_1, u_2)$, are the turn-rate

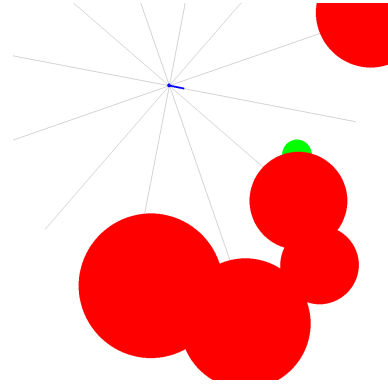


Fig. 2. Dubins4D parallel autonomy simulation environment. Obstacles are randomly placed red circles, the goal is the partially obscured green circle, the agent and its velocity vector is indicated in blue, and the lidar ray casts are light gray. A “reckless” control system navigates the vehicle to the goal without regard to obstacles. We evaluate LRMM, CBF, and HJ-reach methods as “guardian agents” who can temporarily take control of the vehicle to avoid obstacles.

and linear acceleration, respectively. The system is subject to state constraints $v_{\min} \leq v \leq v_{\max}$; control constraints $u_{1,\min} \leq u_1 \leq u_{1,\max}$, $u_{2,\min} \leq u_2 \leq u_{2,\max}$; and bounded disturbance $\|d_x\| \leq d_r$, $\|d_y\| \leq d_r$.

We consider the case where the Dubins4D system in Eqn. 8 is operated in the presences of obstacles and steered by a “reckless driver”, i.e. a controller that guides the system to a goal region but makes no effort to avoid obstacles¹. The objective is to provide a parallel autonomous system—referred to as the “guardian agent”—that judiciously takes control from the driver to prevent collisions with obstacles [34]. It is desired that the guardian agent be minimally-interfering; it should only take control when collisions are imminent and is not responsible for optimizing navigation to the goal.

The environment executes in a *non-blocking* fashion; that is to say that simulation time continues to progress even while agents are computing their next action. This characteristic of the environment is designed to demonstrate the importance of computation time in safety-critical robot applications. The observation spaces is a 17-tuple with simulation time, relative position to goal, absolute heading and speed, and 12 ray-casts equally distributed around the vehicle that sense obstacle locations.

In this experiment we compare three guardian agents: *LRMM*, *CBF*, and *HJ-reach*.

LRMM: The *LRMM agent* is trained with 35,072 samples drawn from 274 randomly generated obstacle configurations that are completely distinct from those configurations used during the evaluations in Table I; i.e. LRMM was *not* allowed to train on the evaluation set. During runtime evaluation, the LRMM agent works by inferring risk from observation of the system’s current state. If the inferred risk is greater than a user-defined threshold ($\hat{\rho}(s) > 0.85$ in our experiments), then the LRMM agent takes active control of the vehicle and

¹For our experiments the “reckless driver” is a control Lyapunov function (CLF) taken from [25]; however, any other controller could be used so long as it generates a non-zero frequency of obstacle collisions.

TABLE I
COMPARATIVE STUDY OF LEARNED RISK METRIC MAPS, CONTROL BARRIER FUNCTIONS, AND HAMILTON-JACOBI REACHABILITY AS PARALLEL AUTONOMY AGENTS IN THE DUBINS4D DRIVING ENVIRONMENT

Agent	Observation Type	Num. Obstacles	Policy Compute Time [ms]	Success Rate [% eps]	Collision Rate [% eps]	Timeout Rate [% eps]	Intervention Rate [% steps]
LRMM	local/partial	5	2.4±5.0	45.3 ± 7.3	29.7±5.6	25.0 ± 4.8	30.7 ± 32.5
CBF	global/privileged	5	52.9 ± 41.2	46.2 ± 5.9	35.1 ± 5.2	18.8 ± 3.8	32.8 ± 39.4
random	none	5	1.4 ± 2.7	27.1 ± 6.9	68.4 ± 5.9	4.5 ± 2.9	~ 50
inactive	none	5	1.0 ± 2.3	44.4 ± 7.4	55.4 ± 7.3	0.1 ± 0.5	0
brakes-only	none	5	0.3 ± 0.7	0.6 ± 0.9	9.0 ± 3.8	90.4 ± 3.7	100
HJ-reach	global/privileged	1	344.0 ± 850.0	54.0 ± 5.2	46.0 ± 5.3	0.0 ± 0.3	< 1.0

applies maximal braking and maximal turning control until the risk estimate drops below the threshold, at which time control is returned to the driver agent².

CBF: The *CBF agent* is based on high-order CBFs described in [25]. CBF-based controllers are most often considered in the context of fully-autonomous systems, but we can easily adapt them to parallel autonomy systems. At each new state that the system encounters, the CBF agent formulates and solves the quadratic program (QP) described in Sec. II. If the CBF constraints are active in the QP solution—i.e. at least one CBF inequality constraint reaches an equality—then the CBF agent overrides the driver and applies the controls resulting from the QP solution. Note that CBFs require an explicit, analytical description of the obstacle space in order to formulate the CBF-based QPs. For our experiments we work around this by providing the CBF agent with privileged, global knowledge of the obstacle space that is not given to the LRMM agent.

HJ reachability: The *HJ-reach agent* is based upon Hamilton-Jacobi reachability [23], [15].

As the Dubins4D driver moves the vehicle through the state space, the HJ-reach agent assess the value function at each state and determines if it is within the unsafe set $\mathcal{V}_{\text{obs}}(t)$. If the system arrives in the unsafe set, then the HJ-reach agent takes control and applies the optimal control to steer the system away from the obstacle set.

Several modifications to the Dubins4D environment were necessary for the HJ-reach agent to be applied. First, like the CBF agent, the HJ-reach agent was granted privileged information about the obstacle space in order to formulate the HJ PDE³. Furthermore, we granted the HJ-reach agent a precomputation phase where the agent is allowed to “peek” at the complete obstacle configuration prior to environment initiation and then given time to solve for the HJ value function over a discretized mesh of the state space. The

²Note that this collision avoidance policy is a heuristic and not derived from the LRMM output. Thus, in this case study, LRMM helps to answer the question of *when* to take control, but not *what* control to apply; i.e. the LRMM is *not prescriptive* of control. See Sec. V for discussion of future work on this topic.

³Unlike CBFs, HJ-reachability does not explicitly require analytical descriptions of the obstacle space and can—in principle—work with an implicit description, e.g. a discretized mesh of the obstacle space. However, our implementation of HJ-reachability was constrained by the Python libraries available for solving the HJ PDE [31] which required explicit obstacle descriptions (e.g. circular regions of known radius and position)

computation of the HJ value function can take up to 2 minutes on the same computer hardware that solves CBF QPs in 50 milliseconds. This is so slow that—without this precompute phase—the entire simulation time window would have elapsed before the HJ-reach agent even had the chance to interact and control the system. Finally, we could only impose a single obstacle for the HJ-reach agent due to limitations of the software library used to implement the HJ solver.

Other Baselines: We also compare against a set of baseline agents—*random*, *inactive*, *brakes-only*—that have simplistic policies that, respectively, intervene with random controls at random times; never intervene on the drivers actions; and always intervene and apply maximum braking at all times. These baselines help bound the possible performance metrics within the Dubins4D environment.

Comparative Results: Table I provides the experimental results from our Dubins4D parallel autonomy case study. A set of 2048 environment configurations (i.e. placement of obstacles, goal, and initial vehicle state) are randomly generated and each agent is evaluated in each of these environment configurations—except the HJ-reach agent which had its own set of 2048 environments due to its limitation to single obstacles. The agents are evaluated based on their policy computation time, success rate (i.e. proportion of trials that end at the goal state), collision rate (trials that end in collision with obstacles), timeout rate (trials that end in neither goal or collision), and intervention rate (i.e. the percent of time steps for which the guardian agent takes control from the driver). Ideally, a guardian agent would have a perfect success rate; a zero collision and timeout rate; and a minimal intervention rate (only takes control at the critical moments to avoid collision and then returns control to the driver); however, such a “perfect” guardian agent is not actually possible in this contrived scenario—see discussion below.

The key insight from Table I is that the LRMM agent performs at least as well—if not better than—CBFs and HJ-reach methods in this real-time, safety-critical, parallel autonomy task. This is in spite of the fact that the LRMM agent only has access to a local, partial observation of the environment whereas the CBF and HJ-reach agents are granted privileged global knowledge of the obstacle space. The success of the LRMM agent is due in large part to its rapid computation time. Needing only to make a feed-

TABLE II
RISK METRIC TRAINING AND TESTING RESULTS ON UNSTRUCTURED ENVIRONMENTS AND HIGH-DIMENSIONAL SYSTEMS

Dynamics	Obstacle Space	State Dims	Control Dims	Observe Dims	Train Samples	Final Train MSE/MAE	Test Samples	Test MSE/MAE
Dubins-3D	Mazes	3	1	8	65536	3.70e-3 1.83e-2	4096	3.38e-2 5.51e-2
Dubins-4D	Circles	4	2	17	35072	2.38e-3 3.71e-2	5260	2.98e-3 3.68e-2
Quadrotor	Walls & Polyhedra	12	4	16	65536	1.69e-4 7.88e-3	8192	4.97e-4 1.13e-2

forward pass through a shallow neural network in order to estimate risk metrics, the LRMM agent can run approximately 20x faster than the CBF agent—which needs to solve quadratic programs at each step—and > 100x faster than the HJ-reach agent—which needs to interpolate a value function and its derivative on a 4-dimensional mesh of the state space. Even though CBFs and HJ reachability can provide theoretical guarantees about control robustness and optimality, these experiments highlight the practical challenges of their implementations in real-time safety-critical applications.

The LRMM and CBF agents both produce roughly the same success rate of 45%. LRMM slightly outperforms CBFs with a 30% collision rate. The Dubins4D environment is randomly initialized such that collisions are unavoidable in some episodes; i.e. the initial state of the vehicle is already within the region of inevitable collision (Eqn. 6) [19]. From the inactive-agent we see that—without intervention from a guardian—the “natural” proportion of episodes ending in collision is roughly 55%. From the brakes-only agent—which always intervenes and immediately brings the vehicle to a stop—we see that an approximate lower bound on collision rate is 9%. LRMM produces a higher rate of episode timeouts—with neither goal nor obstacle being reached—which tend to be caused by the guardian agent bringing the vehicle to a complete stop until the end of the episode. Both LRMM and CBFs have roughly a 30% intervention rate. LRMM and CBF agents significantly out-perform the baseline random agent across all metrics.

Direct comparison with HJ-reach in terms of success, collision, and timeout rate is made difficult by the fact that it was not exposed to the same number of obstacles as other agents. HJ-reach results are included to highlight the long computation time relative to other methods. This is also what leads to its low intervention rate. By the time that the HJ-reach agent has identified that the system has entered an unsafe set and computed the optimal evasion control, the vehicle has already collided with the obstacle.

B. High-Dimensional Systems & Unstructured Environments

As demonstrated in Sec. IV-A, learned risk metric maps can infer a dynamical system’s current risk of failure much more rapidly than existing techniques like CBFs and HJ-reachability. These alternative techniques become even less usable in higher-dimensional systems—where solutions of PDEs over the state space become prohibitively expensive—and unstructured environments—where obstacles cannot be

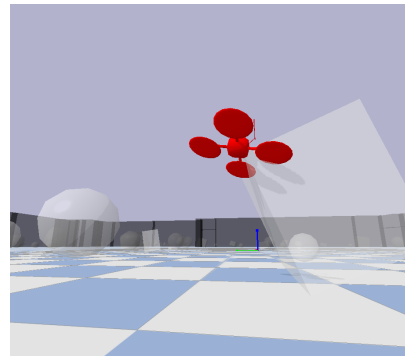


Fig. 3. Visualization of quadrotor in an highly unsafe state (i.e. upside down and near an obstacle). Color of the quadrotor is scaled from green-to-red based on the inferred risk metric which, here, accurately estimates the high risk of the current state. Obstacles and walls made transparent for easier visualization. Rendered in PyBullet [32].

explicitly/analytically described. In this section we show that learned risk metric maps can effectively estimate risk values in such challenging environments. In addition to the Dubins4D environment, we train risk models for a Dubins vehicle in highly-obstructed procedurally-generated mazes (see Fig. 1) and a quadrotor within procedurally-generated rooms with random polyhedron obstacles (see Fig. 3).

Table II gives the risk metric map training and testing results for these three systems. In addition to the MSE training and testing loss we also report mean absolute error (MAE) at the end of training and on the test set. This gives more intuition on how accurately the models estimate risk. We see that all models average a risk estimation error of < 10% of the true risk on the test set.

V. CONCLUSION

In this paper, we present learned risk metric maps (LRMMs) for real-time estimation of coherent risk metrics of high-dimensional dynamical systems operating in unstructured, partially observed environments. We compared the proposed LRMM with other state of the art methods: CBF and HJ-reachability, with results showing advantages of the proposed LRMM’s computation efficiency and general applicability. As noted in Sec. IV-A, a shortcoming of LRMMs is that they estimate failure probabilities, but they do not prescribe the control necessary to avoid failure. We anticipate that future work will remedy this—perhaps through some fusion with control barrier functions—and seek to provide formal guarantees on LRMM accuracy [17].

REFERENCES

- [1] M. Fatemi, T. W. Killian, J. Subramanian, and M. Ghassemi, "Medical dead-ends and learning to identify high-risk states and treatments," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [2] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent measures of risk," *Mathematical finance*, vol. 9, no. 3, pp. 203–228, 1999.
- [3] S. S. Wang, "A class of distortion operators for pricing financial and insurance risks," *Journal of risk and insurance*, pp. 15–36, 2000.
- [4] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," in *Robotics Research*. Springer, 2020, pp. 75–84.
- [5] R. TV, "Maze generator." [Online]. Available: <https://github.com/razimantv/mazegenerator>
- [6] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *ROBOMECH journal*, vol. 1, no. 1, pp. 1–14, 2014.
- [7] A. Wang, X. Huang, A. Jasour, and B. Williams, "Fast risk assessment for autonomous vehicles using learned models of agent futures," *arXiv preprint arXiv:2005.13458*, 2020.
- [8] A. Pierson, W. Schwarting, S. Karaman, and D. Rus, "Navigating congested environments with risk level sets," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5712–5719.
- [9] A. Hakobyan and I. Yang, "Distributionally robust risk map for learning-based motion planning and control: A semidefinite programming approach," *arXiv preprint arXiv:2105.00657*, 2021.
- [10] L. Janson, E. Schmerling, and M. Pavone, "Monte carlo motion planning for robot trajectory optimization under uncertainty," in *Robotics Research*. Springer, 2018, pp. 343–361.
- [11] E. Schmerling and M. Pavone, "Evaluating trajectory collision probability through adaptive importance sampling for safe motion planning," *arXiv preprint arXiv:1609.05399*, 2016.
- [12] A. M. Jasour and B. C. Williams, "Risk contours map for risk bounded motion planning under perception uncertainties." in *Robotics: Science and Systems*, 2019.
- [13] L. Blackmore, M. Ono, A. Bektassov, and B. C. Williams, "A probabilistic particle-control approximation of chance-constrained stochastic predictive control," *IEEE transactions on Robotics*, vol. 26, no. 3, pp. 502–517, 2010.
- [14] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [15] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 2242–2253.
- [16] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [17] T. Lew, L. Janson, R. Bonalli, and M. Pavone, "A simple and efficient sampling-based algorithm for general reachability analysis," in *Learning for Dynamics and Control Conference*. PMLR, 2022, pp. 1086–1099.
- [18] T. Lew and M. Pavone, "Sampling-based reachability analysis: A random set theory approach with adversarial sampling," *arXiv preprint arXiv:2008.10180*, 2020.
- [19] T. Fraichard and H. Asama, "Inevitable collision states—a step towards safer robots?" *Advanced Robotics*, vol. 18, no. 10, pp. 1001–1024, 2004.
- [20] A. Bautin, L. Martínez-Gomez, and T. Fraichard, "Inevitable collision states: a probabilistic perspective," in *2010 IEEE international conference on robotics and automation*. IEEE, 2010, pp. 4022–4027.
- [21] K. Margellos and J. Lygeros, "Hamilton–jacobi formulation for reach-avoid differential games," *IEEE Transactions on automatic control*, vol. 56, no. 8, pp. 1849–1861, 2011.
- [22] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry, "Reach-avoid problems with time-varying dynamics, targets and constraints," in *Proceedings of the 18th international conference on hybrid systems: computation and control*, 2015, pp. 11–20.
- [23] A. Bajcsy, S. Bansal, E. Bronstein, V. Tolani, and C. J. Tomlin, "An efficient reachability-based framework for provably safe autonomous navigation in unknown environments," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 1758–1765.
- [24] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier–value functions for safety-critical control," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 6814–6821.
- [25] W. Xiao and C. Belta, "High-order control barrier functions," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3655–3662, 2022.
- [26] M. J. Kochenderfer, *Decision making under uncertainty: theory and application*. MIT press, 2015.
- [27] R. E. Allen and M. Pavone, "A real-time framework for kinodynamic planning in dynamic environments with application to quadrotor obstacle avoidance," *Robotics and Autonomous Systems*, vol. 115, pp. 174–193, 2019.
- [28] I. A. Sucas, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.
- [30] M. S. Andersen, J. Dahl, L. Vandenberghe *et al.*, "Cvxopt: A python package for convex optimization," *abel. ee. ucla. edu/cvxopt*, vol. 88, 2013. [Online]. Available: <https://cvxopt.org/>
- [31] M. Bui, G. Giovanis, M. Chen, and A. Shriraman, "Optimizeddp: An efficient, user-friendly library for optimal control and dynamic programming," *arXiv preprint arXiv:2204.05520*, 2022.
- [32] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.
- [33] R. Soklaski, J. Goodwin, O. Brown, M. Yee, and J. Matterer, "Tools and practices for responsible ai engineering," *arXiv preprint arXiv:2201.05647*, 2022.
- [34] A. Amini, G. Rosman, S. Karaman, and D. Rus, "Variational end-to-end navigation and localization," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8958–8964.