

A Quality Diversity Approach to Automatically Generating Human-Robot Interaction Scenarios in Shared Autonomy

Matthew C. Fontaine and Stefanos Nikolaidis
Department of Computer Science
University of Southern California
Los Angeles, CA, USA
mfontain@usc.edu, nikolaid@usc.edu

Abstract—The growth of scale and complexity of interactions between humans and robots highlights the need for new computational methods to automatically evaluate novel algorithms and applications. Exploring diverse scenarios of humans and robots interacting in simulation can improve understanding of the robotic system and avoid potentially costly failures in real-world settings. We formulate this problem as a quality diversity (QD) problem, where the goal is to discover diverse failure scenarios by simultaneously exploring both environments and human actions. We focus on the shared autonomy domain, where the robot attempts to infer the goal of a human operator, and adopt the QD algorithm MAP-Elites to generate scenarios for two published algorithms in this domain: shared autonomy via hindsight optimization and linear policy blending. Some of the generated scenarios confirm previous theoretical findings, while others are surprising and bring about a new understanding of state-of-the-art implementations. Our experiments show that MAP-Elites outperforms Monte-Carlo simulation and optimization based methods in effectively searching the scenario space, highlighting its promise for automatic evaluation of algorithms in human-robot interaction.

I. INTRODUCTION

We present a method for automatically generating human-robot interaction (HRI) scenarios in shared autonomy. Consider as an example a manipulation task, where a user provides inputs to a robotic manipulator through a joystick, guiding the robot towards a desired goal, e.g., grasping a bottle on the table. The robot does not know the goal of the user in advance, but infers their desired goal in real-time by observing their inputs and assisting them by moving autonomously towards that goal. Performance of the algorithm is assessed by how fast the robot reaches the goal. However, different environments and human behaviors could cause the robot to fail, by picking the wrong object or colliding with obstacles.

Typically, such algorithms are evaluated with human subject experiments [81]. While these experiments are fundamental in exploring and evaluating human-robot interactions and they can lead to exciting and unpredictable behaviors, they are often limited in the number of environments and human actions that they can cover. Testing an algorithm in simulation with a *diverse* range of scenarios can improve understanding of the system, inform the experimental setup of real-world studies,




Fig. 1: An example archive of solutions returned by the quality diversity algorithm MAP-Elites. The solutions in red indicate scenarios where the robot fails to reach the desired user goal in a simulated shared autonomy manipulation task. The scenarios vary with the environment (y-axis: distance between the two candidate goals) and human inputs (x-axis: variation from optimal path).

and help avoid potentially costly failures “in the wild.”

One approach is to simulate agent behaviors by repeatedly sampling from models of human behavior and interaction protocols [80]. While this approach will show the *expected* behavior of the system given the pre-specified models, it is unlikely to reveal failure cases that are not captured by the models or are in the tails of the sampling distribution. Exhaustive search of human actions and environments is also computationally prohibitive given the continuous, high-dimensional space of all possible environments and human action sequences.

Another approach is to formulate this as an optimization

problem, where the goal is to find adversarial environments and human behaviors. But we are typically not interested in the maximally adversarial scenario, which is the single, global optimum of our optimization objective, since these scenarios are both easy to find and unlikely to occur in the real-world, e.g., the human moving the joystick of an assistive robotic arm consistently in the wrong direction.

Instead, we are interested in answering questions of the form: how noisy can the human input be before the algorithm breaks? Or, in the aforementioned example task, how far apart do two candidate goals have to be for the robot to disambiguate the human intent?

Our work makes the following contributions:¹

1. We propose formulating the problem of generating human-robot interaction scenarios as a *quality diversity* (QD) problem, where the goal is not to find a single, optimal solution, but a collection of high-quality solutions, in our case failure scenarios of the tested algorithm, across a range of measurable criteria, such as noise in human inputs and distance between objects.

2. We adopt the QD algorithm MAP-Elites, originally presented in [14, 62], for the problem of scenario generation. Focusing on the shared autonomy domain, where a robotic manipulator attempts to infer the user’s goal based on their inputs, we show that MAP-Elites outperforms two baselines: standard Monte Carlo simulation (random search), where we uniformly sample the scenario parameters, and CMA-ES [34], a state-of-the-art derivative-free optimization algorithm, in finding diverse scenarios that minimize the performance of the tested algorithm. We select to test the algorithm “shared autonomy with hindsight optimization” [43], since it has been widely used and we have found it to perform robustly in a range of different environments and tasks. Additionally, in hindsight optimization, inference and planning are tightly coupled, which makes testing particularly challenging; simply testing each individual component is not sufficient to reveal how the algorithm will perform.

3. We show that Monte Carlo simulation does not perform well because of *behavior space distortion*: sampling directly from the space of environments and human actions covers only a small region in the space of measurable aspects (behavioral characteristics). For example, uniformly sampling object locations (scenario parameters) results in a non-uniform distribution of their distances (behavioral characteristic) with a very small variance near the mean. On the other hand, MAP-Elites focuses on exploring the space of the behavioral characteristics by retaining an archive of high-performing solutions in that space and perturbing existing solutions with small variations. Therefore, MAP-Elites performs a type of simultaneous search guided by the behavioral characteristics, where solutions in the archive are used to generate future candidate solutions [62].

4. We analyze the failure cases and we show that they result from specific aspects of the implementation of the tested

algorithm, rather than being artifacts of the simulation environment. We use the same approach to contrast the performance of hindsight optimization with that of linear policy blending [19] and generate a diverse range of scenarios that confirm previous theoretical findings [82]. The generated scenarios transfer to the real-world; we reproduce some of the automatically discovered scenarios on a real robot with human inputs. While some of the scenarios are expected, e.g., the robot approaches the wrong goal if the human provides very noisy inputs, others are surprising, e.g, the robot never reaches the desired goal even for a nearly optimal user if the two objects are aligned in column formation in front of the robot (Fig. 1)!

QD algorithms treat the algorithm being tested as a “black box”, without any knowledge of its implementation, which makes them applicable to multiple domains. Overall, we are excited about the potential of QD to facilitate understanding of complex HRI systems, opening up a number of scientific challenges and opportunities to be explored in the future.

II. PROBLEM STATEMENT

Given a shared autonomy system where a robot interacts with a human, our goal is to generate scenarios that minimize performance of the system, while ensuring that the generated scenarios cover a range of prespecified measures.

We let R be a single robot interacting with a single human H . We assume a function G_H that generates human inputs, a function G_E that generates an environment, and an HRI algorithm G_R that generates actions for the robot. The human input generator is parameterized by $\theta \in \mathbb{R}^{n_\theta}$, where n_θ is the dimensionality of the parameter space, while the environment generator is parameterized by $\phi \in \mathbb{R}^{n_\phi}$. We define a *scenario* as the tuple (θ, ϕ) .

In shared autonomy, $G_E(\phi)$ generates an initial environment (and robot) state x_E . The human observes x_E and provides inputs to the system $u_H = G_H(x_E, \theta)$ through some type of interface. The robot observes x_E and the human input u_H and takes an action $u_R = G_R(x_E, u_H)$. The state changes with dynamics: $\dot{x}_E = h(x_E, u_R)$. H and R interact for a time horizon T , or until they reach a final state $x_f \in X_E$.

To evaluate a scenario, we assume a function $f(x_E^{0..T}, u_R^{0..T}, u_H^{0..T}) \rightarrow \mathbb{R}$ that maps the state and action history to a real number. We call this an *assessment* function, which measures the performance of the robotic system. We also assume M user-defined functions, $b_i(x_E^{0..T}, u_R^{0..T}, u_H^{0..T}) \rightarrow \mathbb{R}$, $i \in [M]$. These functions measure aspects of generated scenarios that should vary, e.g., noise in human inputs or distance between obstacles. We call these functions *behavior characteristics* (BCs), which induce a Cartesian space called a *behavior space*.

Given the parameterization of the environment and human input generators, we can map a value assignment of the parameters (θ, ϕ) to a state and action history $(x_E^{0..T}, u_R^{0..T}, u_H^{0..T})$ and therefore to an assessment $f(\theta, \phi)$ and a set of BCs $b(\theta, \phi)$. We assume that the behavior space is partitioned into N cells, which form an *archive* of scenarios, and we let (θ_i, ϕ_i) be the parameters of the scenario occupying cell $i \in [N]$.

¹We include an overview video: <https://youtu.be/9P3qomydMWk>

The objective of our scenario generator is to fill in as many cells of the archive as possible with scenarios of high assessment f :²

$$\mathcal{M}(\theta_1, \phi_1, \dots, \theta_N, \phi_N) = \max \sum_{i=1}^N f(\theta_i, \phi_i) \quad (1)$$

III. BACKGROUND

Automatic Scenario Generation Automatically generating scenarios is a long standing problem in human training [38], with the core challenge being the generation of *realistic* scenarios [59]. Previous work [89] has shown that optimization methods can be applied to generate scenarios by maximizing a scenario quality metric.

Scenario generation has been applied extensively to evaluating autonomous vehicles [6, 65, 2, 76, 29, 77]. Contrary to model-checking and formal methods [12, 68], which require a model describing the system’s performance such as a finite-state machine [60] or process algebras [68], black-box approaches do not require access to a model. Most relevant are black-box falsification methods [15, 88, 47, 20] that attempt to find an input trace that minimizes the performance of the tested system. Rather than searching for a single global optimum [15, 16, 77], or attempting to maximize coverage of the space of scenario parameters [88] or of performance boundary regions [65], we propose a quality diversity approach where we optimize an archive formed by a set of behavioral characteristics, with a focus on the shared autonomy domain. This allows us to *simultaneously* search for human-robot interaction scenarios that minimize the performance of the system over a range of measurable criteria, e.g., over a range of variation in human inputs and distance between goal objects.

Finally, scenario generation is closely related to the problem of generating video game levels in procedural content generation (PCG) [36, 78]. An approach gaining popularity is procedural content generation through quality diversity (PCG-QD) [32], which leverages QD algorithms to drive the search for interesting and diverse content.

Quality Diversity and MAP-Elites. QD algorithms differ from pure optimization methods, in that they do not attempt to find a single optimal solution, but a collection of good solutions that differ across specified dimensions of interest. For example, QD algorithms have generated video game levels of varying number of enemies or tile distributions [48, 25], and objects of varying shape complexity and grasp difficulty [61].

MAP-Elites [62, 14] is a popular QD algorithm that searches along a set of explicitly defined attributes called *behavior characteristics* (BCs), which induce a Cartesian space called a *behavior space*. The behavior space is tessellated into uniformly spaced grid cells. In each cell, the algorithm maintains the highest performing solution, which is called as *elite*. The

collection of elites returned by the algorithm forms an *archive* of solutions.

MAP-Elites populates the archive by first randomly sampling a population of solutions, and then selecting the elites – which are the top performing solutions in each cell of the behavior space – at random and perturbing them with small variations. The objective of the algorithm is two-fold: maximize the number of filled cells (coverage) and maximize the quality of the elite in each cell. Recent algorithms have focused on how the behavior space is tessellated [79, 24], as well as how each elite is perturbed [83]. Recent work [26] has also shown that, when the objective function and behavior characteristics are first-order differentiable, MAP-Elites via a Gradient Arborescence (MEGA) can result in significant improvements in search efficiency.

By retaining an archive of high-performing solutions and perturbing existing solutions with small variations, MAP-Elites simultaneously optimizes every region of the archive, using existing solutions as “stepping stones” to find new solutions. Previous work has shown that MAP-Elites variants [63] and surrogate models [28] outperform independent single-objective constrained optimizations for each cell with CMA-ES, with the same total budget of evaluations.

Coverage-Driven Testing in HRI. Previous work [4, 5] explored test generation in human-robot interaction using Coverage-Driven Verification (CDV), emulating techniques used in functional verification of hardware designs. Human action sequences were randomly generated in advance and with a model-based generator which modeled the interaction with Probabilistic-Timed Automata. Instead, we focus on online scenario generation by searching over a set of scenario parameters; the generator itself is agnostic to the underlying HRI algorithm and human model. Previous work [3] has also used Q-learning to generate plans for an agent in order to maximize coverage. Our focus is both on coverage and quality of generating scenarios, with respect to a prespecified set of behavioral characteristics that we want to cover. In contrast to previous studies that simulate human actions, *we jointly search for environments and human/agent behaviors*.

Shared Autonomy. Shared autonomy (also: shared control, assistive teleoperation) combines human teleoperation of a robot with intelligent robotic assistance. The method has been applied in the control of robotic arms [42, 19, 67, 64, 37, 31, 40, 44, 57, 72, 71], the flight of UAVs [74, 30, 53], and robot-assisted surgery [54, 75]. Shared autonomy has been implemented through a variety of interfaces, such as whole body motions [19], natural language [18], laser pointers [84], brain-computer interfaces [64], body-machine interfaces [41], and eye gaze [10, 42]. A shared autonomy system first predicts the human’s goal, often through machine learning methods trained from human demonstrations [35, 50, 86], and then provides assistance, which often involves blending the user’s input with the robot assistance to achieve the predicted goal [19, 21, 49]. Assistance can provide task-dependent guidance [1], manipulation of objects [44], or mode switches [37].

Shared Autonomy via Hindsight Optimization. In shared

²We note that the assessment function could be any performance metric of interest, such as time to completion or minimum robot’s distance to obstacles. Additionally, while in this work we focus on minimizing performance, we could instead search for scenarios that maximize performance, or that achieve performance that matches a desired value. We leave this for future work.

autonomy via hindsight optimization [43] assistance blends user input and robot control based on the confidence of the robot’s goal prediction. The problem is formulated as a Partially Observable Markov Decision Process (POMDP), wherein the user’s goal is a latent variable. The system models the user as an approximately optimal stochastic controller, which provides inputs so that the robot reaches the goal as fast as possible. The system treats the user’s inputs as observations to update a distribution over the user’s goal, and assists the user by minimizing the expected cost to go – estimated using the distance to goal – for that distribution. Since solving a POMDP exactly is intractable, the system uses the hindsight optimization (QMDP) approximation [56]. The system was shown to achieve significant improvements in efficiency of manipulation tasks in an object-grasping task [43] and more recently in a feeding task [42]. We empirically found this algorithm to perform robustly in a range of different environments, which motivates a systematic approach for testing. We refer to this algorithm simply as *hindsight optimization*.

IV. SCENARIO GENERATION WITH MAP-ELITES

Algorithm 1 shows the MAP-Elites algorithm from [62, 14], adapted for scenario generation. The algorithm takes as input a function G_H parameterized by θ that generates human inputs, a function G_E parameterized by ϕ that generates environments, and an HRI algorithm G_R that generates actions for the robot. The algorithm searches for scenarios θ, ϕ of high assessment values f that fill in the archive \mathcal{P} .

For each scenario (θ, ϕ) , MAP-Elites instantiates the generator functions G_H and G_E . For instance, ϕ could be a vector of objects positions, and θ could be a vector of waypoints representing a trajectory of human inputs, or parameters of a human policy.

When a scenario is generated, MAP-Elites executes the scenario in a simulated environment and estimates the assessment function f and the BCs b . MAP-Elites then updates the archive if (1) the cell corresponding to the BCs $\mathcal{X}[b]$ is empty, or (2) the existing scenario (elite) in $\mathcal{X}[b]$ has a smaller assessment function (lower quality) than the new scenario. This allows populating the archive to maximize coverage as well as improving the quality of existing scenarios.

For the first N_{init} iterations, the algorithm generates scenarios θ, ϕ by randomly sampling from the parameter space. These sampled parameters seed the archive with an initial set of scenarios. After the first N_{init} iterations, MAP-Elites selects a scenario uniformly at random from the archive and perturbs it with a small variation. This allows for better exploration of the archive, compared to random search, as we show in section VI-E.

We note that, while our experiments focus on the shared autonomy domain, the proposed scenario generation method is general and can be applied to multiple HRI domains.

V. GENERATING SCENARIOS IN SHARED AUTONOMY

We focus on a shared autonomy manipulation task, where a human user teleoperates a robotic manipulator through a

Algorithm 1 Scenario Generation with MAP-Elites

Input: Human input generator G_H , environment generator G_E , HRI algorithm G_R , variations $\sigma_\theta, \sigma_\phi$
Initialize: Scenarios in archive $\mathcal{X} \leftarrow \emptyset$, assessments $\mathcal{F} \leftarrow \emptyset$
for $t = 1, \dots, N$ **do**
 if $t < N_{init}$ **then**
 Generate scenario $\theta, \phi = \text{random_generation}()$
 else
 Select elite $\theta', \phi' = \text{random_selection}(\mathcal{X})$
 Sample $\theta \sim N(\theta', \sigma_\theta)$
 Sample $\phi \sim N(\phi', \sigma_\phi)$
 end if
 Instantiate $G_H^\theta = G_H(\theta)$
 Instantiate $G_E^\phi = G_E(\phi)$
 Compute $f = \text{assessment}(G_H^\theta, G_E^\phi, G_R)$
 Compute $b = \text{behaviors}(G_H^\theta, G_E^\phi, G_R)$
 if $\mathcal{F}[b] = \emptyset$ or $\mathcal{F}[b] < f$ **then**
 Update archive $\mathcal{X}[b] \leftarrow (\theta, \phi), \mathcal{F}[b] \leftarrow f$
 end if
end for

joystick interface. The robot runs a hindsight optimization shared autonomy algorithm [43], which uses the user’s input to infer the object the user wants the robot to grasp, and assists the user by moving autonomously towards that goal.

A. Scenario Parameters

Following the specification of section II, we define a human input generator G_H parameterized by θ and an environment generator G_E parameterized by ϕ .

Environment Generator: The environment generator G_E takes as input the 2D positions g_i of n goal objects (bottles), so that $\phi = (g^1, \dots, g^n)$, and places them on top of a table. We specify the range of the coordinates $g_x \in [0, 0.25]$ (in meters), $g_y \in [0, 0.2]$ so that the goals are always reachable by the robot’s end-effector. We position the robotic manipulator to face the objects (Fig. 3).

Human Input Generator: Our pilot studies with the shared autonomy system have shown that user inputs are typically not of constant magnitude. Instead, inputs spike when users wish to “correct” the robot’s path, and decrease in magnitude afterwards when the robot takes over.

Therefore, we specify the human input generator G_H , so that it generates a set of equidistant waypoints in Cartesian space forming a straight line that starts from the initial position of the robot’s end-effector and ends at the desired goal of the user. At each timestep, the generator takes as input the current state of the robot (and the environment) x_E , and provides a translational velocity command u_H for the robot’s end-effector towards the next waypoint, proportional to the distance to that waypoint.

We allow for noise in the waypoints by adding a disturbance $d \in [-0.05, 0.05]$ for each of the intermediate waypoints in the horizontal direction (x -axis in Fig. 3). We selected $m = 5$

intermediate waypoints, and specified the human input parameter θ as a vector of disturbances, so that $\theta = (d_1, \dots, d_5)$. We note that this is only one way, out of many, of simulating the human inputs.

HRI Algorithm and Simulation Environment: We use the publicly available implementation of the hindsight optimization algorithm [52], which runs on the OpenRAVE [17] simulation environment. Experiments were conducted with a Gen2 Lightweight manipulator. For each goal object we assume one target grasp location, on the side of the object that is facing the robot.

B. Assessment Function.

The assessment function f represents the quality of a scenario. We evaluate a scenario by simulating it until the robot’s end-effector reaches the user’s goal, or when the maximum time (10 s) has elapsed. We use as an assessment function time to completion, where longer times represent higher scenario quality, since we wish to discover scenarios that *minimize* performance.

C. Behavioral Characteristics.

We wish to generate scenarios that show the limits of the shared autonomy system: how noisy can the human be without the system failing to reach the desired goal? How does distance between candidate goals affect the system’s performance? Intuitively, noisier human inputs and smaller distances between goals would make the inference of the user’s goal harder and thus make the system more likely to fail.

These dimensions of interest are the behavioral characteristics (BC) b : attributes that we wish to obtain coverage for. We explore the following BCs:

Distance Between Goals: How far apart the human goal is from other candidate goals in a scenario plays an important role in disambiguating the human intent when the robot runs the hindsight optimization algorithm. The reason is that the implementation of the algorithm models the human user as minimizing a cost function proportional to the distance between the robot and the desired goal. The framework then infers the user’s goal by using the user inputs as observations; the more unambiguous the user input, the more accurate the inference of the system. Therefore, we expect that the further away the human goal object g_H is from the nearest goal g_N , the better the system will perform. We define this BC as:

$$BC1 = \|g_H - g_N\|_2 \quad (2)$$

Given the range of the goal coordinates, the range of this BC is [0, 0.32]. In practice, there will be always a minimum distance between two goal objects because of collisions, but this does not affect our search, since we can ignore cases where the objects collide. We partitioned this behavior space to 25 uniformly spaced intervals.

Human Variation: We expect noise in the human inputs to affect the robot’s inference of the user’s goal and thus the system’s performance. We capture variation from the optimal

path using the root sum of the squares of the disturbances d_i applied to the m intermediate waypoints.

$$BC2 = \sqrt{\sum_{i=1}^m d_i^2} \quad (3)$$

A value of 0 indicates a straight line to the goal. Since we have $d_i \in [-0.05, 0.05]$ (section V-A), the range of this BC is [0, 0.11]. We partitioned this behavior space to 100 uniformly spaced intervals.

Human Rationality: If we interpret the user’s actions using a bounded rationality model [8, 22], we can explain deviations from the optimal trajectory of human inputs as a result of how “rational” or “irrational” the user is.³

Formally, we let x_R be the 3D position of the robot’s end-effector and u_H be the velocity controlled by the user in Cartesian space. We model the user as following Boltzmann policy $\pi_H \mapsto P(u_H|x_R, g_H, \beta)$, where β is the rationality coefficient – also interpreted as the expertise [44]– of the user and Q_{g_H} is the value function from x_R to the goal g_H .

$$P(u_H|x_R, g_H, \beta) \propto e^{-\beta Q_{g_H}(x_R, u_H)} \quad (4)$$

Let $Q_{g_H} = -\|u_H\|_2 - \|x_R + u_H - g_H\|_2$ [22], so that the user minimizes the distance to the goal. Observe that if $\beta \rightarrow \infty$, the human is rational, providing velocities exactly in the direction to the goal. If $\beta \rightarrow 0$, the human is random, choosing actions uniformly.

We can estimate the user’s rationality, given their inputs, with Bayesian inference [22]:

$$P(\beta|x_R, g_H, u_H) \propto P(u_H|x_R, g_H, \beta)P(\beta) \quad (5)$$

Since the human inputs change at each waypoint (section V-A), we perform $m + 1$ updates, at the starting position and at each intermediate waypoint, on a finite set of discrete values β . Following previous work [44], we set the rationality range $\beta \in [0, 1000]$. We then choose as behavioral characteristic the value with the maximum a posteriori probability at the end of the task:

$$BC3 = \operatorname{argmax} P(\beta|x_R^{0:T}, g_H, u_H^{0:T}) \quad (6)$$

We partitioned the space to 101 uniformly spaced intervals.

VI. EXPERIMENTS

We compare different search algorithms in their ability to find diverse and high-quality scenarios in different behavior spaces.

³We note that we use the human rationality model as one way, out of many, to *interpret* human inputs and not to as a way to *generate* inputs. Human inputs can be generated with any generator model. In this paper, we generate human inputs with the deterministic model described in section V-A. We discuss extensions to stochastic human models in section VIII.

	BC1 & BC3, 2 goals		BC1 & BC2, 2 goals		BC1 & BC2, 3 goals	
Algorithm	Coverage	QD-Score	Coverage	QD-Score	Coverage	QD-score
Random	22.3%	3464	48.4%	7782	41.9%	7586
CMA-ES	24.8%	4540	38.9%	7422	34.5%	7265
MAP-Elites	62.8%	10128	63.0%	11216	57.4%	11204

TABLE I: Results: Percentage of cells covered (coverage) and QD-Score after 10,000 evaluations, averaged over 5 trials.

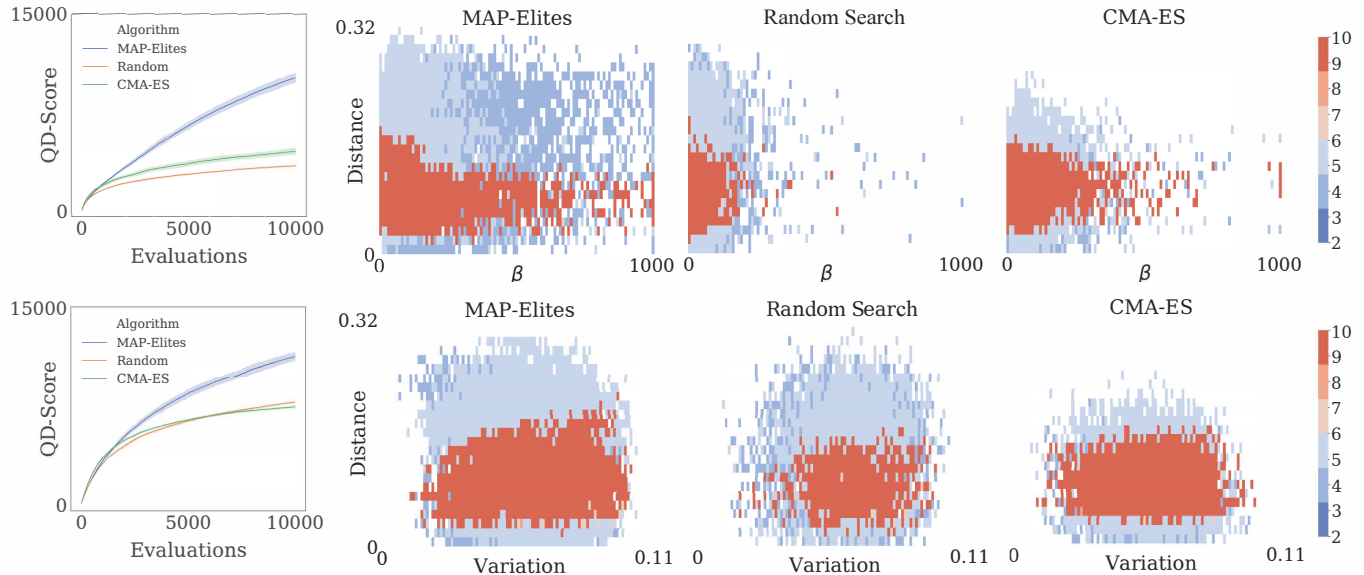


Fig. 2: QD-Scores over evaluations (generated scenarios) and example archives returned by the three algorithms for the first two behavior spaces of Table I. The colors of the cells in the archives represent time to task completion in seconds.

A. Independent Variables

The experiment has two independent variables, the *behavior space* and the *search algorithm*.

Behavior Space: (1) Distance between $n = 2$ goal objects (BC1) and human rationality (BC3), (2) Distance between $n = 2$ goal objects (BC1) and human variation (BC2), and (3) Distance between human goal and nearest goal for $n = 3$ goals (BC1) and human variation (BC2).⁴

Search Algorithm: We evaluate three different search methods: MAP-Elites, CMA-ES and random search. The Covariance Matrix Adaptation Evolution Strategy (CMA-ES) is one of the most competitive derivative-free optimizers for single-objective optimization of continuous spaces (see [34, 33]) and it is commonly used for falsification of cyber-physical systems [16, 87]. In random search we use Monte Carlo simulation where scenario parameters are sampled uniformly within their prespecified ranges.

We implemented a multi-processing system on an AMD Ryzen Threadripper 64-core (128 threads) processor, as a master search node and multiple worker nodes running separate OpenRAVE processes in parallel, which enables simultaneous

⁴We note that the behavior spaces can be more than two-dimensional, e.g. we could specify a space with all three BCs. We include only 2D spaces since they are easier to visualize and inspect.

evaluation of many scenarios. Random search and MAP-Elites run asynchronously on the master search node, while CMA-ES synchronizes before each covariance matrix update. We generated 10,000 scenarios per trial, and ran 45 trials, 5 for each algorithm and behavior space. One trial parallelized into 100 threads lasted approximately 20 minutes.

B. Algorithm Tuning

MAP-Elites first samples uniformly the space of scenario parameters θ, ϕ within their prespecified ranges for an initial population of 100 scenarios (Algorithm 1). The algorithm then randomly perturbs the elites (scenarios from the archive) with Gaussian noise scaled by a σ parameter. The two scenario parameters, position of goal objects ϕ and human waypoints θ , are on different scales, thus we specified a different σ for each: $\sigma_\phi = 0.01, \sigma_\theta = 0.005$.

To generate the scenarios for random search, we uniformly sample scenario parameters within their prespecified ranges, a method identical to generating the initial population of MAP-Elites.

For CMA-ES, we selected a population of $\lambda = 12$ following the recommended setting from [34]. To encourage exploration, we used the bi-population variant of CMA-ES with restart rules [7, 33], where the population doubles after each restart, and we selected a large step size, $\sigma = 0.05$. Since the two

search parameters are in different scales, we initialized the diagonal elements of the covariance matrix C , so that $c_{ii} = 1.0, i \in [2n]$ and $c_{ii} = 0.5, i \in \{2n + 1, \dots, 2n + m\}$, with $2n$ and m the dimensionality of the goal object and human input parameter spaces respectively.

Both CMA-ES and MAP-Elites may sample scenario parameters that do not fall inside their prespecified ranges. Following recent empirical results on bound constraint handling [9], we adopted a resampling strategy, where new scenarios are resampled until they fall within the prespecified range.

C. Measures

We wish to measure both the diversity and the quality of scenarios returned by each algorithm. These are combined by the QD-Score metric [70], which is defined as the sum of f values of all elites in the archive (Eq. 1 in section II). Empty cells have 0 f value. Therefore, QD-score is positively affected by both the coverage of the archive (the number of occupied cells in the archive divided by the total number of cells) and the assessment of the occupied cells. Similarly to previous work [27], we compute the QD-Score of CMA-ES and random search for comparison purposes by calculating the behavioral characteristics for each scenario and populating a pseudo-archive. We also include the coverage score as an additional metric of diversity.

D. Hypothesis

We hypothesize that MAP-Elites will result in larger QD-Score and coverage than both CMA-ES and random search.

Previous work [24, 27] has shown that behavior spaces are typically distorted: uniformly sampling the search parameter space results in samples concentrated in small areas of the behavior space. Therefore, we expect random search to have small coverage of the behavior space. Additionally, since random search ignores the assessment function f , we expect the quality of the found scenarios in the archive to be low.

CMA-ES moves with a single large population that has global optimization properties. Therefore, we expect it to concentrate in regions of high-quality scenarios, rather than explore the archive. On the other hand, MAP-Elites both expands the archive and maximizes the quality of the scenarios within each cell.

E. Analysis

Table I summarizes the performance of the three algorithms, for each of the three behavior spaces. We conducted a two-way ANOVA to examine the effect of the behavior space and the search algorithm on the QD-Score and coverage. There was a statistically significant interaction between the search algorithm and the behavior space for both QD-Score ($F(4, 36) = 62.39, p < 0.001$) and coverage ($F(4, 36) = 77.92, p < 0.001$). Simple main effects analysis with Bonferroni correction showed that MAP-Elites outperformed CMA-ES and random search in both QD-Score and coverage ($p < 0.001$ in all comparisons). This result supports our hypothesis.

Fig. 2 shows the improvement in the QD-Score over time and one example archive from each algorithm for the first two behavior spaces. MAP-Elites visibly finds more cells and of higher quality (red color), illustrating its ability to cover larger areas of the archive with high-quality scenarios. As expected, CMA-ES concentrates in regions of high-quality scenarios but has small coverage.

Random search covers a smaller area of the archive, compared to MAP-Elites, because of the *behavior space distortion*, shown in Fig. 4. Even though the search parameters are sampled uniformly, scenarios are concentrated on the left side of the archive specified by the human rationality and distance between goals BCs (Fig. 4-top). This occurs because if any of the sampled waypoints deviates from the optimal path, low values of rationality become more likely. In the human variation and distance between goals BCs, the distribution of scenarios generated by random search is concentrated in a small region near the center (Fig. 4-bottom). This is expected, since the two BCs are Euclidean norms of random vectors (see [85]). On the other hand, MAP-Elites selects elite scenarios from the archive and perturbs them with Gaussian noise, instead of uniformly sampling the scenario parameters, resulting in larger coverage.

F. Interpreting the Archives

In the generated archives (Fig. 2), each cell contains an elite, which is the scenario that achieved the maximum assessment value (time to complete the task) for that cell. It is important to confirm that *the timeouts (red cells in the archive) occur because of the implementation of the tested algorithm (hindsight optimization), rather than being artifacts of the simulation environment.*

Therefore, we replay the elites in different regions of the archives to explain the system’s performance. We focus on the first two behavior spaces in Fig. 2 using the archives generated with MAP-Elites, since MAP-Elites had the largest QD-Score and coverage.

We observe that if the distance between goals is large and the human is nearly optimal, the robot performs the task efficiently. This is shown by the blue color in the top-right of the first behavior space (distance and human rationality β). We observe the same for large distance and small variation in the second behavior space.

We then explore different types of scenarios where the robot fails to reach the user’s goal by the maximum time (10s), indicated by the red cells in the archives. When human variation is large (or equivalently rationality is low), the human may provide inputs that guide the robot towards the wrong goal. Since the robot updates a probability distribution over goals based on the user’s input [43] and the robot assumes that the user minimizes their distance to their desired goal, noisy inputs may result in assigning a higher probability to the wrong goal and the robot will move towards that goal instead. Fig. 3(left) shows the execution trace of one elite where this occurs. Fig. 1 shows the position of this elite in the archive.




Fig. 3: (Left) The robot fails to reach the user’s goal g_H because of the large deviation in human inputs from the optimal path. The waypoints of the human inputs are indicated with green color. (Center) We show for comparison how the robot would act if human deviation was 0 (optimal human). (Right) The robot fails to reach the user’s goal g_H (bottle furthest away from the robot), even though the human provides a near optimal input trajectory.

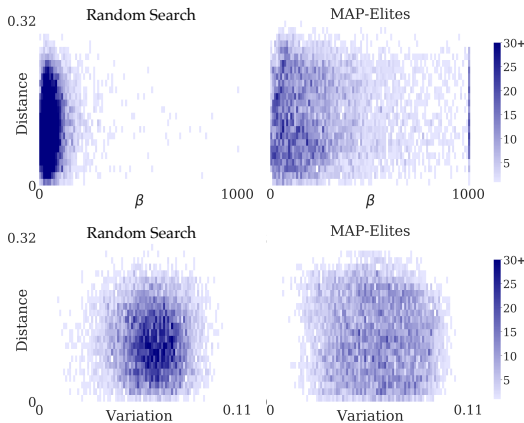


Fig. 4: Distribution of cells explored for random search and MAP-Elites. The cell colors represent frequency counts.

Fig. 3(center) shows how the robot would reach the desired goal if the human had behaved optimally, instead.

What is surprising, however, is that the robot does not reach the user’s goal even in parts of the behavior space where human variation is nearly 0 (or equivalently rationality is very high), that is when the human provides a nearly optimal input trajectory! Fig. 3(right) reveals a case, where the two goal objects are aligned one closely behind the other. The robot approaches the first object, on the way towards the second object, and stops there.

What is interesting in both scenarios is that the robot gets “stuck” at the wrong goal, even when the simulated user continues providing inputs to their desired goal! Inspection of the publicly available implementation [52] of the algorithm shows that this results from the combination of two factors: the robot’s cost function and the human inputs.

Cost Function. The cost function that the robot minimizes is specified as a constant cost when the robot is far away from the goal and as a smaller linear cost when the robot is near the target [42] (distance to target is smaller than a threshold). This makes the cost of the goal object near the robot significantly lower than the cost of the other goal objects, which results in the probability mass of the goal prediction to concentrate on that goal. While this can help the user align the end-effector with the object (see [42]), it can also lead to

incorrect inference, if the robot approaches the wrong goal on its way towards the correct goal or because of noisy human input. We confirmed that removing the linear term from the cost function results in the robot reaching the right goal in both examples.

Human Inputs. The hindsight optimization implementation minimizes a cost function specified as the sum of two quadratic terms, the expected cost-to-go to assist for a distribution over goals, and a term penalizing disagreement with the user’s input. The first-order approximation of the value function leads to an interpretation of the final robot action $u_R = u_R^A + u_R^u$ as the sum of two independent velocity commands, an “autonomous” action u_R^A towards the distribution over goals and an action that follows the user input u_R^u , as if the user was directly teleoperating the robot (see [42]).

We have simulated the human inputs, so that they provide a translational velocity command towards the next waypoint, proportional to the distance of the robot’s end-effector to that waypoint (section V-A). This results in a term u_R^u of small magnitude when the end-effector is close to one of the waypoints. If at the same time the robot has high confidence on one of the goals, u_R^A will point in the direction of that goal and it will cancel out any term u_R^u that attempts to move the robot in the opposite direction.

We confirmed that, if the user instead applied the maximum possible input towards their desired goal, the robot would get “unstuck,” so a real user would always be able to eventually reach their desired goal. However, this requires effort from the user who would need to “fight” the robot. Overall, the archive reveals limitations that depend on *how the goal objects are aligned in the environment, the direction and magnitude of user inputs, and the cost function used by the implementation of the hindsight optimization algorithm.*

VII. COMPARING ALGORITHMS

Given the effectiveness of quality diversity in automatically generating a diverse range of test scenarios, we can also use it to understand differences in performance between algorithms. In this work, we compare the performance of hindsight optimization [43] with linear policy blending [19]. We describe the experiment in Appendix A. We found that policy blending resulted in collisions, even for a nearly optimal human, in cases where human and robot inputs pointed towards opposite




Fig. 5: Scenarios where the policy blending algorithm results in collision with an obstacle, approximated by a sphere. (Left) While the human and robot trajectories are each collision-free, blending the two results to collision when they point towards opposite sides of the obstacle. (Right) Blending with a very noisy human input results in collision.




Fig. 6: We reproduced the generated scenarios in the real world with actual joystick inputs.

sides of an obstacle (Fig. 5). This confirms previous theoretical findings [82] of unsafe behavior that linear blending has in the presence of obstacles. On the other hand, hindsight optimization avoids collisions, since it uses the human inputs as observations and the robot’s motion is determined only by the robot’s policy.

VIII. DISCUSSION

Experimental Findings. We found that failure scenarios for hindsight optimization occur when the two goals are close to each other and the human inputs are noisy, or when one goal is in front of the other. In the latter case, failure occurs even if the human input is nearly optimal in minimizing the distance to the desired goal. In both cases, the robot becomes overconfident about the wrong goal and gets “stuck” there.

An important factor is the linear decrease of the cost in the vicinity of the goal objects. When specifying the cost function, it would be prudent to make the distance threshold for the linear decrease proportional to the distance between the goal objects, rather than setting it to an absolute value.

Other potential measures to avoid the system’s overconfidence towards the wrong goal are: (1) including the Shannon entropy with respect to all the goals in the cost function [44] to penalize actions that result in very high confidence to one specific goal; (2) assigning a non-zero probability that the user changes their mind throughout the task and switches goals [66, 39]. It would be interesting to investigate the effect of more “conservative” assistance on subjective and objective metrics of the robot’s performance.

Finally, while linear policy blending naturally gives more

control to the user and it is preferred by users in simple tasks [42], we empirically verified that the algorithm can generate unsafe trajectories, even if the individual human and robot inputs are safe.

To show that the presented scenarios can occur in deployed systems, we reproduce them in the real world with actual inputs through a joystick interface (Fig. 6).⁵

Stochasticity in Scenarios. In our experiments the generated scenarios are deterministic. One may wish to simulate scenarios where there is stochasticity in the robot’s decision making or in the environment dynamics. A designer may also wish to test the system’s performance under a stochastic human model, e.g., when human inputs are generated by a stochastic noisily rational human.

The most common approach in evolutionary optimization of noisy domains is *explicit averaging*, where we run multiple trials of the same scenario and then retain an aggregate measure of the assessment estimate, e.g. we compute the average to estimate the expected assessment $\mathbb{E}[f(\theta, \phi)]$ [73, 45]. We can follow the same process to estimate the behavior characteristics [46]. To improve the efficiency of the estimation, previous work has also employed implicit averaging methods, e.g., where the assessment of a scenario (θ, ϕ) is estimated by taking the assessments of previously evaluated scenarios in the neighborhood of θ, ϕ into account. Previous work also includes adaptive sampling techniques, where the number of trials increases over time as the quality of the solutions in the archive improves [46]. A recent variant of MAP-Elites (Deep-Grid MAP-Elites) which updates a subpopulation of solutions for each cell in the behavior space has shown significant benefits in sample efficiency [23]. We leave these exciting directions for future work.

Limitations. An important challenge is how to effectively characterize the behavior spaces. While we have assumed bounded behavior spaces, the rationality coefficient does not meet this assumption, which resulted in elites accumulating in the upper bound of the rationality in the archive (Fig. 2). Adapting the boundaries of the space dynamically based on the distribution of generated scenarios [24] could improve coverage in this case.

While distance between objects indeed played a role, our experiments showed the unexpected and surprising edge case where the two objects are in column formation and the human is nearly optimal. An interesting follow-up experiment would be to specify as a BC some metric of object alignment in column formation and investigate further the effect of this variable. In general, a practitioner can test the system with an initial design of BCs, observe the failure cases, create new BCs from newly observed insights and test the system further.

We focused on how to effectively search the generative space of scenarios, but not on the generation methods themselves. Realism is an important future consideration, both in generating environments and human inputs. In human training,

⁵We show different generated scenarios reproduced in the real world here: <https://youtu.be/2-JCO3dUHsA>

realism can be measured through a modified Turing test designed to require humans to distinguish generated scenarios from human authored ones [58]. Alternatively, we could run a user study where we place objects in the same locations as our failure scenarios and observe whether participants perform similar actions that cause failures.

Implications. Finding failure scenarios of HRI algorithms will play a critical role in the adoption of these systems. We proposed quality diversity as an approach for automatically generating scenarios that assess the performance of HRI algorithms in the shared autonomy domain and we illustrated a path for future work. While real-world studies are essential in evaluating complex HRI systems, automatic scenario generation can facilitate understanding and tuning of the tested algorithms, as well as provide insights on the experimental design of real-world studies, whose findings can in turn inform the designer for testing the system further. We are excited about applications of quality diversity algorithms as test oracles in verification systems [51, 69], as well as in other domains where deployed robotic systems face a diverse range of interaction scenarios.

IX. ACKNOWLEDGEMENTS

We thank Tapomayukh Bhattacharjee, David Hsu, Shen Li, Dylan Losey, Dorsa Sadigh, Rosario Scalice, Julian Togelius and our anonymous RSS reviewers for their feedback on early versions of this work.

REFERENCES

- [1] Daniel Aarno, Staffan Ekvall, and Danica Kragic. 2005. Adaptive virtual fixtures for machine-assisted teleoperation tasks. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*. IEEE, 1139–1144.
- [2] Y. Abeysirigoonawardena, F. Shkurti, and G. Dudek. 2019. Generating Adversarial Driving Scenarios in High-Fidelity Simulators. In *2019 International Conference on Robotics and Automation (ICRA)*. 8271–8277. <https://doi.org/10.1109/ICRA.2019.8793740>
- [3] Dejanira Araiza-Illan, Anthony G Pipe, and Kerstin Eder. 2016. Intelligent agent-based stimulation for testing robotic software in human-robot interactions. In *Proceedings of the 3rd Workshop on Model-Driven Robot Software Engineering*. 9–16.
- [4] Dejanira Araiza-Illan, David Western, Anthony Pipe, and Kerstin Eder. 2015. Coverage-Driven Verification—. In *Haifa Verification Conference*. Springer, 69–84.
- [5] Dejanira Araiza-Illan, David Western, Anthony G Pipe, and Kerstin Eder. 2016. Systematic and realistic testing in simulation of control code for robots in collaborative human-robot interactions. In *Annual Conference Towards Autonomous Robotic Systems*. Springer, 20–32.
- [6] James Arnold and Rob Alexander. 2013. Testing Autonomous Robot Control Software Using Procedural Content Generation. In *Proceedings of the 32Nd International Conference on Computer Safety, Reliability, and Security - Volume 8153 (SAFECOMP 2013)*. Springer-Verlag New York, Inc., New York, NY, USA, 33–44. https://doi.org/10.1007/978-3-642-40793-2_4
- [7] Anne Auger and Nikolaus Hansen. 2005. A restart CMA evolution strategy with increasing population size. In *2005 IEEE congress on evolutionary computation*, Vol. 2. IEEE, 1769–1776.
- [8] Chris L Baker, Joshua B Tenenbaum, and Rebecca R Saxe. 2007. Goal inference as inverse planning. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 29.
- [9] Rafał Biedrzycki. 2020. Handling bound constraints in CMA-ES: An experimental study. *Swarm and Evolutionary Computation* 52 (2020), 100627.
- [10] Zeungnam Bien, Myung-Jin Chung, Pyung-Hun Chang, Dong-Soo Kwon, Dae-Jin Kim, Jeong-Su Han, Jae-Hean Kim, Do-Hyung Kim, Hyung-Soon Park, Sang-Hoon Kang, et al. 2004. Integration of a rehabilitation robotic system (KARES II) with human-friendly man-machine interaction units. *Autonomous robots* 16, 2 (2004), 165–191.
- [11] Tom Carlson and Yiannis Demiris. 2012. Collaborative control for a robotic wheelchair: evaluation of performance, attention, and workload. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 42, 3 (2012), 876–888.
- [12] Jiyoung Choi. 2013. Model checking for decision making behaviour of heterogeneous multi-agent autonomous system. (2013).
- [13] Jacob W Crandall and Michael A Goodrich. 2002. Characterizing efficiency of human robot interaction: A case study of shared-control teleoperation. In *IEEE/RSJ international conference on intelligent robots and systems*, Vol. 2. IEEE, 1290–1295.
- [14] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503.
- [15] Jyotirmoy Deshmukh, Marko Horvat, Xiaoqing Jin, Rupa Majumdar, and Vinayak S Prabhu. 2017. Testing cyber-physical systems through bayesian optimization. *ACM Transactions on Embedded Computing Systems (TECS)* 16, 5s (2017), 1–18.
- [16] Jyotirmoy Deshmukh, Xiaoqing Jin, James Kapinski, and Oded Maler. 2015. Stochastic local search for falsification of hybrid systems. In *International Symposium on Automated Technology for Verification and Analysis*. Springer, 500–517.
- [17] Rosen Diankov and James Kuffner. 2008. Openrave: A planning architecture for autonomous robotics. *Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34* 79 (2008).
- [18] Finale Doshi and Nicholas Roy. 2007. Efficient model learning for dialog management. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*. ACM, 65–72.
- [19] A.D. Dragan and S.S. Srinivasa. 2012. Formalizing

- Assistive Teleoperation. In *Proc. Robotics: Science and Systems Conference*.
- [20] Tommaso Dreossi, Daniel J Fremont, Shromona Ghosh, Edward Kim, Hadi Ravanbakhsh, Marcell Vazquez-Chanlatte, and Sanjit A Seshia. 2019. Verifai: A toolkit for the formal design and analysis of artificial intelligence-based systems. In *International Conference on Computer Aided Verification*. Springer, 432–442.
- [21] Andrew H Fagg, Michael Rosenstein, Robert Platt, and Roderic A Grupen. 2004. Extracting user intent in mixed initiative teleoperator control. In *Proceedings of the American Institute of Aeronautics and Astronautics Intelligent Systems Technical Conference*.
- [22] Jaime F Fisac, Andrea Bajcsy, Sylvia L Herbert, David Fridovich-Keil, Steven Wang, Claire J Tomlin, and Anca D Dragan. 2018. Probabilistically safe robot planning with confidence-based human predictions. *arXiv preprint arXiv:1806.00109* (2018).
- [23] Manon Flageat and Antoine Cully. 2020. Fast and stable MAP-Elites in noisy domains using deep grids. In *Artificial Life Conference Proceedings*. MIT Press, 273–282.
- [24] Matthew C. Fontaine, Scott Lee, L. B. Soros, Fernando de Mesentier Silva, Julian Togelius, and Amy K. Hoover. 2019. Mapping Hearthstone Deck Spaces through MAP-Elites with Sliding Boundaries. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '19)*. ACM, New York, NY, USA, 161–169. <https://doi.org/10.1145/3321707.3321794>
- [25] Matthew C Fontaine, Ruilin Liu, Julian Togelius, Amy K Hoover, and Stefanos Nikolaidis. 2021. Illuminating mario scenes in the latent space of a generative adversarial network. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2021).
- [26] Matthew C Fontaine and Stefanos Nikolaidis. 2021. Differentiable Quality Diversity. *arXiv e-prints* (2021), arXiv–2106.
- [27] Matthew C Fontaine, Julian Togelius, Stefanos Nikolaidis, and Amy K Hoover. 2020. Covariance Matrix Adaptation for the Rapid Illumination of Behavior Space. *Proceedings of the Genetic and Evolutionary Computation Conference* (2020).
- [28] Adam Gaier, Alexander Asteroth, and Jean-Baptiste Mouret. 2017. Data-efficient exploration, optimization, and modeling of diverse designs through surrogate-assisted illumination. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 99–106.
- [29] Alessio Gambi, Marc Mueller, and Gordon Fraser. 2019. Automatically Testing Self-driving Cars with Search-based Procedural Content Generation. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA 2019)*. ACM, New York, NY, USA, 318–328. <https://doi.org/10.1145/3293882.3330566>
- [30] Jeremy H Gillula, Gabriel M Hoffmann, Haomiao Huang, Michael P Vitus, and Claire J Tomlin. 2011. Applications of hybrid reachability analysis to robotic aerial vehicles. *The International Journal of Robotics Research* 30, 3 (2011), 335–354.
- [31] Deepak Gopinath, Siddarth Jain, and Brenna D Argall. 2016. Human-in-the-loop optimization of shared autonomy in assistive robotics. *IEEE Robotics and Automation Letters* 2, 1 (2016), 247–254.
- [32] Daniele Gravina, Ahmed Khalifa, Antonios Liapis, Julian Togelius, and Georgios N Yannakakis. 2019. Procedural content generation through quality diversity. In *2019 IEEE Conference on Games (CoG)*. IEEE, 1–8.
- [33] Nikolaus Hansen. 2009. Benchmarking a BI-population CMA-ES on the BBOB-2009 function testbed. In *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers*. 2389–2396.
- [34] Nikolaus Hansen. 2016. The CMA evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772* (2016).
- [35] Kris Hauser. 2013. Recognition, prediction, and planning for assisted teleoperation of freeform tasks. *Autonomous Robots* 35, 4 (2013), 241–254.
- [36] Mark Hendriks, Sebastiaan Meijer, Joeri Van Der Velden, and Alexandru Iosup. 2013. Procedural content generation for games: A survey. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 9, 1 (2013), 1–22.
- [37] Laura V Herlant, Rachel M Holladay, and Siddhartha S Srinivasa. 2016. Assistive teleoperation of robot arms via automatic time-optimal mode switching. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*. IEEE Press, 35–42.
- [38] Ronald C Hofer, Edward Ramirez, and Scott H Smith. 1998. Automated Scenario Generation Environment. (1998).
- [39] Siddarth Jain and Brenna Argall. 2018. Recursive Bayesian human intent recognition in shared-control robotics. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3905–3912.
- [40] Siddarth Jain and Brenna Argall. 2019. Probabilistic human intent recognition for shared autonomy in assistive robotics. *ACM Transactions on Human-Robot Interaction (THRI)* 9, 1 (2019), 1–23.
- [41] Siddarth Jain, Ali Farshchiansadegh, Alexander Broad, Farnaz Abdollahi, Ferdinando Mussa-Ivaldi, and Brenna Argall. 2015. Assistive robotic manipulation through shared autonomy and a body-machine interface. In *2015 IEEE international conference on rehabilitation robotics (ICORR)*. IEEE, 526–531.
- [42] Shervin Javdani, Henny Admoni, Stefania Pellegrinelli, Siddhartha S Srinivasa, and J Andrew Bagnell. 2018. Shared autonomy via hindsight optimization for teleoperation and teaming. *The International Journal of Robotics Research* (2018), 0278364918776060.
- [43] S. Javdani, S. S. Srinivasa, and J.A. Bagnell. 2015. Shared Autonomy via Hindsight Optimization. In *Proc.*

Robotics: Science and Systems Conference.

- [44] Hong Jun Jeon, Dylan P Losey, and Dorsa Sadigh. 2020. Shared Autonomy with Learned Latent Actions. *arXiv preprint arXiv:2005.03210* (2020).
- [45] Yaochu Jin and Jürgen Branke. 2005. Evolutionary optimization in uncertain environments—a survey. *IEEE Transactions on evolutionary computation* 9, 3 (2005), 303–317.
- [46] Niels Justesen, Sebastian Risi, and Jean-Baptiste Mouret. 2019. MAP-Elites for noisy domains by adaptive sampling. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 121–122.
- [47] James Kapinski, Jyotirmoy V Deshmukh, Xiaoqing Jin, Hisahiro Ito, and Ken Butts. 2016. Simulation-based approaches for verification of embedded control systems: An overview of traditional and advanced modeling, testing, and verification techniques. *IEEE Control Systems Magazine* 36, 6 (2016), 45–64.
- [48] Ahmed Khalifa, Michael Cerny Green, Gabriella Barros, and Julian Togelius. 2019. Intentional computational level design. In *Proceedings of The Genetic and Evolutionary Computation Conference*. 796–803.
- [49] Jonathan Kofman, Xianghai Wu, Timothy J Luu, and Siddharth Verma. 2005. Teleoperation of a robot manipulator using a vision-based human-robot interface. *IEEE Transactions on Industrial Electronics* 52, 5 (2005), 1206–1219.
- [50] Hema S Koppula and Ashutosh Saxena. 2016. Anticipating human activities using object affordances for reactive robotic response. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 1 (2016), 14–29.
- [51] Hadas Kress-Gazit, Kerstin Eder, Guy Hoffman, Henny Admoni, Brenna Argall, Ruediger Ehlers, Christoffer Heckman, Nils Jansen, Ross Knepper, Jan Křetínský, et al. 2020. Formalizing and Guaranteeing* Human-Robot Interaction. *arXiv preprint arXiv:2006.16732* (2020).
- [52] Personal Robotics Lab. 2017. https://github.com/personalrobotics/ada_assistance_policy/.
- [53] Thanh Mung Lam, Harmen Wigert Boschloo, Max Mulder, and Marinus M Van Paassen. 2009. Artificial force field for haptic feedback in UAV teleoperation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 39, 6 (2009), 1316–1330.
- [54] Ming Li and Allison M Okamura. 2003. Recognition of operator motions for real-time assistance using virtual fixtures. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2003. HAPTICS 2003. Proceedings. 11th Symposium on*. IEEE, 125–131.
- [55] Qinan Li, Weidong Chen, and Jingchuan Wang. 2011. Dynamic shared control for human-wheelchair cooperation. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, 4278–4283.
- [56] Michael L Littman, Anthony R Cassandra, and Leslie Pack Kaelbling. 1995. Learning policies for partially observable environments: Scaling up. In *Machine Learning Proceedings 1995*. Elsevier, 362–370.
- [57] Dylan P Losey, Krishnan Srinivasan, Ajay Mandlekar, Animesh Garg, and Dorsa Sadigh. 2019. Controlling assistive robots with learned latent actions. *arXiv preprint arXiv:1909.09674* (2019).
- [58] Glenn A. Martin. 2012. *Automatic Scenario Generation using Procedural Modeling Techniques*. Ph.D. Dissertation. University of Central Florida.
- [59] Glenn A. Martin, Charles E. Hughes, Sae Schatz, and Denise Nicholson. 2010. The Use of Functional L-systems for Scenario Generation in Serious Games. In *Proceedings of the 2010 Workshop on Procedural Content Generation in Games (PCGames '10)*. ACM, New York, NY, USA, Article 6, 5 pages. <https://doi.org/10.1145/1814256.1814262>
- [60] Karl Meinke and Peter Nycander. 2015. Learning-based testing of distributed microservice architectures: Correctness and fault injection. In *SEFM 2015 Collocated Workshops*. Springer, 3–10.
- [61] Douglas Morrison, Peter Corke, and Jurgen Leitner. 2020. EGAD! an Evolved Grasping Analysis Dataset for diversity and reproducibility in robotic manipulation. *IEEE Robotics and Automation Letters* (2020).
- [62] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [63] Jean-Baptiste Mouret and Glenn Maguire. 2020. Quality diversity for multi-task optimization. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. 121–129.
- [64] Katharina Muelling, Arun Venkatraman, Jean-Sebastien Valois, John E. Downey, Jeffrey Weiss, Shervin Javdani, Martial Hebert, Andrew B. Schwartz, Jennifer L. Collinger, and J. Andrew Bagnell. 2017. Autonomy Infused Teleoperation with Application to Brain Computer Interface Controlled Manipulation. *Auton. Robots* 41, 6 (Aug. 2017), 1401–1422.
- [65] Galen E. Mullins, Paul G. Stankiewicz, R. Chad Hawthorne, and Satyandra K. Gupta. 2018. Adaptive generation of challenging scenarios for testing and evaluation of autonomous vehicles. *Journal of Systems and Software* 137 (2018), 197 – 215. <https://doi.org/10.1016/j.jss.2017.10.031>
- [66] Stefanos Nikolaidis, David Hsu, and Siddhartha Srinivasa. 2017. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research* 36, 5-7 (2017), 618–634.
- [67] S. Nikolaidis, Zhu. Y., D. Hsu, and S.S. Srinivasa. 2017. Human-Robot Mutual Adaptation in Shared Autonomy. In *Proc. Human Robot Interaction Conference*.
- [68] Matthew O’Brien, Ronald C Arkin, Dagan Harrington, Damian Lyons, and Shu Jiang. 2014. Automatic verification of autonomous robot missions. In *International Conference on Simulation, Modeling, and Programming for Autonomous Robots*. Springer, 462–473.
- [69] David Porfirio, Allison Sauppé, Aws Albarghouthi, and

- Bilge Mutlu. 2018. Authoring and verifying human-robot interactions. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 75–86.
- [70] Justin K Pugh, Lisa B Soros, Paul A Szerlip, and Kenneth O Stanley. 2015. Confronting the challenge of quality diversity. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. 967–974.
- [71] Daniel Rakita, Bilge Mutlu, Michael Gleicher, and Laura M Hiatt. 2018. Shared dynamic curves: A shared-control telemanipulation method for motor task training. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 23–31.
- [72] Daniel Rakita, Bilge Mutlu, Michael Gleicher, and Laura M Hiatt. 2019. Shared control-based bimanual robot manipulation. *Science Robotics* 4, 30 (2019).
- [73] Pratyusha Rakshit, Amit Konar, and Swagatam Das. 2017. Noisy evolutionary optimization algorithms—a comprehensive survey. *Swarm and Evolutionary Computation* 33 (2017), 18–45.
- [74] Siddharth Reddy, Sergey Levine, and Anca Dragan. 2018. Shared Autonomy via Deep Reinforcement Learning. *arXiv preprint arXiv:1802.01744* (2018).
- [75] Jing Ren, Rajni V Patel, Kenneth A McIsaac, Gerard Guiraudon, and Terry M Peters. 2008. Dynamic 3-D virtual fixtures for minimally invasive beating heart procedures. *IEEE transactions on medical imaging* 27, 8 (2008), 1061–1070.
- [76] E. Rocklage, H. Kraft, A. Karatas, and J. Seewig. 2017. Automated scenario generation for regression testing of autonomous vehicles. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. 476–483. <https://doi.org/10.1109/ITSC.2017.8317919>
- [77] Dorsa Sadigh, S Shankar Sastry, and Sanjit A Seshia. 2019. Verifying robustness of human-aware autonomous cars. *IFAC-PapersOnLine* 51, 34 (2019), 131–138.
- [78] Noor Shaker, Julian Togelius, and Mark J. Nelson. 2016. *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*. Springer.
- [79] Davy Smith, Laurissa Tokarchuk, and Geraint Wiggins. 2016. Rapid phenotypic landscape exploration through hierarchical spatial partitioning. In *International conference on parallel problem solving from nature*. Springer, 911–920.
- [80] Aaron Steinfeld, Odest Chadwicke Jenkins, and Brian Scassellati. 2009. The oz of wizard: simulating the human for interaction research. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*. 101–108.
- [81] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. 2016. Computational human-robot interaction. *Foundations and Trends in Robotics* 4, 2-3 (2016), 105–223.
- [82] Pete Trautman. 2015. Assistive planning in complex, dynamic environments: a probabilistic approach. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 3072–3078.
- [83] Vassilis Vassiliades and Jean-Baptiste Mouret. 2018. Discovering the Elite Hypervolume by Leveraging Interspecies Correlation. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 149–156.
- [84] Eduardo Veras, Karan Khokar, Redwan Alqasemi, and Rajiv Dubey. 2009. Scaled telerobotic control of a manipulator in real time with laser assistance for ADL tasks. In *Mechatronics and its Applications, 2009. ISMA'09. 6th International Symposium on*. IEEE, 1–6.
- [85] Roman Vershynin. 2018. *Random Vectors in High Dimensions, page 3869*. Cambridge University Press.
- [86] Zhikun Wang, Katharina Mülling, Marc Peter Deisenroth, Heni Ben Amor, David Vogt, Bernhard Schölkopf, and Jan Peters. 2013. Probabilistic movement modeling for intention inference in human-robot interaction. *The International Journal of Robotics Research* 32, 7 (2013), 841–858.
- [87] Zhenya Zhang, Gidon Ernst, Ichiro Hasuo, and Sean Sedwards. 2018. Time-staging enhancement of hybrid system falsification. In *2018 IEEE Workshop on Monitoring and Testing of Cyber-Physical Systems (MT-CPS)*. IEEE, 3–4.
- [88] Qianchuan Zhao, Bruce H Krogh, and Paul Hubbard. 2003. Generating test inputs for embedded control systems. *IEEE Control Systems Magazine* 23, 4 (2003), 49–57.
- [89] Alexander Zook, Stephen Lee-Urban, Mark O Riedl, Heather K Holden, Robert A Sottilare, and Keith W Brawner. 2012. Automated scenario generation: toward tailored and optimized military training in virtual environments. In *Proceedings of the international conference on the foundations of digital games*. 164–171.

APPENDIX A COMPARING HINDSIGHT OPTIMIZATION WITH POLICY BLENDING

Given the effectiveness of quality diversity in automatically generating a diverse range of test scenarios, we can also use it to understand differences in performance between algorithms. Using the shared autonomy domain as example, we compare hindsight optimization [43] with policy blending [19].

Policy Blending. In shared autonomy, policy blending is a widely used alternative to POMDP-based methods in shared autonomy, since it provides an intuitive interface for combining robotic assistance with user inputs [11, 55, 31, 64, 19, 40]. In this approach, the robot’s and user’s actions are treated as two independent sources. The final robot motion is determined by an arbitration function that regulates the two inputs, based on the robot’s confidence about the user’s goal. A special case of policy blending is *linear blending*. Theoretical analysis of linear blending [82] has shown that it can lead to unsafe behavior in the presence of obstacles.

We want to empirically confirm these theoretical findings and assess how hindsight optimization and linear policy blending perform in the presence of obstacles. We thus created a

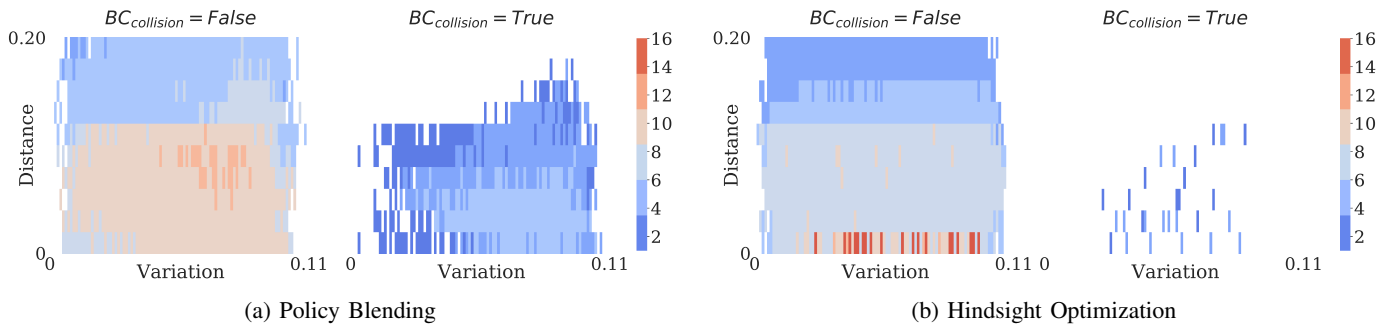


Fig. 7: Archives generated with MAP-Elites for the policy blending and hindsight optimization algorithms. For each algorithm, we present the 3D behavior space as two 2D spaces, one where there was no collision ($BC_{collision} = False$), and one where a collision occurred ($BC_{collision} = True$). The colors of the cells in the archives represent duration of the task in seconds.

new environment with one goal object and a second object that acts as an obstacle. Since there is only one unique goal object, there will be no failure cases arising from incorrect inference as in section V.

In linear blending, robot and human inputs are weighted by the arbitration coefficients α and $1 - \alpha$ respectively. Previous work [19] considers two types: an “aggressive” blending, where the robot takes over ($\alpha = 1.0$) if confidence is above a threshold ($conf \geq 0.4$), and a “timid” blending where α linearly increases from 0 to 0.6 for $0.4 \leq conf \leq 0.8$. When we have only one goal, confidence is always 1.0 and “aggressive” blending and hindsight optimization are identical. Therefore, we consider only the “timid” mode which sets α to 0.6.

Obstacle Modeling. While the hindsight optimization implementation [52] models the user as minimizing the Euclidean distance between the robot’s end-effector and their goal, this is no longer applicable in the presence of obstacles, since a direct path to the goal may collide with the obstacle. Exact computation of the value function is also infeasible, given that the state and action spaces are continuous and the robot acts in real-time. We simplify the computation by modeling the obstacle as a sphere and, if a direct path to the goal intersects with the sphere, we find the intersecting points and approximate the value function as the length of the shortest path to the goal that wraps around the sphere. We do the same for the human waypoints (before adding any disturbances), so that an optimal human will follow the shortest collision-free path to the goal. For robustness, if the end-effector touches the surface of the sphere, it receives an additional velocity command in the direction away from the center of the sphere, similarly to a potential field [13]. We found the sphere approximation simplification adequate for detecting collisions with the hindsight optimization and policy blending algorithms.

Experiment. We fix the position of the goal object and the obstacle in the y-axis, so that the obstacle o is always between the robot and the goal, and we search for the coordinates ($\phi = (g_x, o_x)$) in the x-axis (Fig. 5). The goal object has three target grasp poses. Identically to section V, we search

for the disturbances to the human inputs ($\theta = (d_1, \dots, d_5)$). We specify three BCs: the distance between the two objects in the x-axis, the human variation – identically to section V – and a binary BC, $BC_{collision} \in \{True, False\}$ indicating whether the robot’s end-effector has collided with the obstacle. As before, we set the assessment function f to the time to completion, with a maximum limit of 15s. The task terminates if the robot reaches the goal or if it collides with the obstacle. We generate 20,000 scenarios with one run of MAP-Elites to test the two algorithms.

Analysis. Fig. 7 shows the archives generated for the policy blending and the hindsight optimization algorithms. For each algorithm, we present two 2D archives, one with $BC_{collision} = True$ and one with $BC_{collision} = False$. We compare the archives of the two algorithms in terms of coverage and quality of the scenarios.

Coverage. The coverage was 47% for hindsight optimization, compared to 69% for policy blending. We observe that the archive with $BC_{collision} = True$ is heavily populated in policy blending (Fig. 7a). The archive shows collisions even for a nearly optimal human; while both the human and robot inputs were collision free, sometimes they pointed towards opposite sides of the obstacle, and blending the two resulted in collision (Fig. 5(left)). This result matches the theoretical predictions [82] of unsafe trajectories in linear blending.

The $BC_{collision} = True$ archive in Fig. 7a also shows that when the horizontal distance exceeds a threshold, the robot can reach the goal with a nearly straight path, and collision occurs only when the human inputs are very noisy (Fig. 5(right)).

On the other hand, hindsight optimization resulted in a nearly empty archive when $BC_{collision} = True$. In other words, there were very few scenarios where a collision has occurred. The reason is that hindsight optimization uses the human inputs as observations, and the robot’s motion is determined only by the robot’s policy; the few sparse collisions are artifacts of the OpenRAVE environment, occurring because of infrequent lags in sending velocity commands to the robot.

Scenario Assessment. When $BC_{collision} = False$, the average time to completion was 7.95s for policy blending and 6.70s for hindsight optimization, indicating that policy blending took

more time to finish the task. This result matches previous human subjects experiments [42], where policy blending took longer than hindsight optimization, and it was caused by two factors: (1) Often the human and robot inputs would point to opposite directions, resulting in velocities of small magnitudes when blended. (2) The human inputs were proportional to the distance to the next waypoint, spiking at the waypoints and decreasing until the next waypoint (section V-A). This resulted in parts of the trajectory where the human inputs were of smaller magnitude than the robot's, so blending the two gave smaller velocity commands than treating the human inputs as observations. The red bars in hindsight optimization (Fig. 7b) indicate timeouts; when the goal object was exactly behind the obstacle, the robot sometimes oscillated between taking a grasping position and moving away from the sphere.

Overall, the generated scenarios show two potential limitations of linear policy blending, compared to hindsight optimization: unsafe behavior in the presence of obstacles, and delays in task completion, when human and robot inputs point to opposite directions.