

Generative adversarial imitation learning for robot swarms: Learning from human demonstrations and trained policies

Mattes Kraus¹ and Jonas Kuckling^{1,2,3}

Abstract—In imitation learning, robots are supposed to learn from demonstrations of the desired behavior. Most of the work in imitation learning for swarm robotics provides the demonstrations as rollouts of an existing policy. In this work, we provide a framework based on generative adversarial imitation learning that aims to learn collective behaviors from human demonstrations. Our framework is evaluated across six different missions, learning both from manual demonstrations and demonstrations derived from a PPO-trained policy. Results show that the imitation learning process is able to learn qualitatively meaningful behaviors that perform similarly well as the provided demonstrations. Additionally, we deploy the learned policies on a swarm of TurtleBot 4 robots in real-robot experiments. The exhibited behaviors preserved their visually recognizable character and their performance is comparable to the one achieved in simulation.

I. INTRODUCTION

A robot swarm is a decentralized, homogeneous multi-robot systems, in which each robots acts autonomously based only on the information that is locally available to it [1], [2]. These properties lend themselves to the creation of multi-robot systems that exhibit some desirable properties, such as scalability, flexibility and robustness [2]. While robot swarms have shown to be able to tackle a wide range of missions [1], [3], [4], designing control software for any particular mission remains a key challenge in swarm robotics [5], [6]. This is due to the fact that the swarm needs to be programmed on the individual level, but the desired collective behavior is emergent from the numerous and often unpredictable interactions of the robots. Manual design approaches therefore often rely on trial-and-error, and therefore depend on the skill of the designer [7].

Automatic methods transform the design problem into an optimization problem [7]. Instead of finding suitable control software to address a specific mission, an optimization algorithm searches for the instance of control software that maximizes some (mission-specific) performance measure. It is generally assumed that the instance of control software that

maximizes the mission-specific performance measure also solves the mission in a satisfactory way. However, designing these performance measures is itself a challenging problem. If they only measure the desired outcomes, they might not provide enough information during early stages of the design process [8], [9], [10]. Contrarily, if the designer also encodes “how” the behavior should be achieved, the performance measure becomes prone to reward hacking [11]. As a result, high-scoring behaviors might not correspond to desired ones.

Imitation learning approaches offer an alternative way to specify desired behaviors. Instead of providing an explicit performance measure, the designer provides *demonstrations* of the desired behavior. The imitation learning algorithm will then derive a policy that exhibits the same distribution of *features* as the demonstrations.

In this work, we propose an imitation learning framework for robot swarms that uses generative adversarial imitation learning (GAIL) [12]. We focus on swarm-level features to capture emergent collective dynamics, rather than focusing on the behavior of individual agents. Demonstrations are provided by human operators that can control the robot swarm through a custom-developed demonstration tool. To evaluate the impact of human-provided demonstrations, we also consider imitation learning from demonstrations that are derived from policies trained with proximal policy optimization (PPO) [13]. The framework is further validated through real-robot experiments, deploying the learned policies on a real swarm of TurtleBot 4 robots.

II. RELATED WORK

Works that consider imitation learning in swarm robotics usually employ forms of behavior cloning, explicit feature matching, inverse reinforcement learning, or generative adversarial imitation learning. In behavior cloning, the demonstrations are used to collect sets of state-action (or observation-action) pairs. The problem is then cast as a supervised learning problem, i.e., the policy is trained to predict the corresponding actions to a states given as input to it [14], [15], [16]. In the case of explicit feature matching, the demonstrations are used to compute a set of features (on the agent- or swarm-level) describing the desired behavior. The imitation learning process then aims to explicitly minimize the difference between the features in the demonstrations and the features of the generated behaviors [17], [18]. In inverse reinforcement learning approaches, the difference penalty for feature distributions is not explicitly given. Instead, the features are used to learn a reward function, which might assign different importance to some features [19], [20], [21].

¹ University of Konstanz, Department of Computer and Information Science, Konstanz, Germany

² Cluster of Excellence “Centre for the Advanced Study of Collective Behavior”, University of Konstanz

³ Zukunftskolleg, University of Konstanz

*JK acknowledges funding from the Carl-Zeiss-Stiftung and from the Zukunftskolleg. The research was supported by the Cluster of Excellence “Centre for the Advanced Study of Collective Behavior”, under Germany’s Excellence Strategy – EXC 2117 – 422037984.

The experiments were designed by MK and JK and performed by MK. Real robot experiments were performed by MK and JK. The paper was written and edited by MK and JK, all authors read and approved the final version. The research was directed by JK.

jonas.kuckling@uni-konstanz.de

In generative adversarial imitation learning approaches, no explicit objective function exists. Instead, a discriminator and generator (the policy) are trained in parallel [12]. The policy generates state-action pairs and the discriminator tries to decide if any given state-action pair was part of the demonstration or is generated. The discriminator is trained to decrease the classification error and the policy is rewarded if it generates state-action pairs that the discriminator incorrectly classifies as originating from the demonstration. A first work of generative adversarial imitation learning in multi-robot systems is Turing learning [22]. Li et al. use co-evolution to evolve one discriminator network and generator model per imitating agent. The discriminator only has access to the current linear and angular speed of the agent. Wu et al. use state-action pairs to train the discriminator using OpenAI-ES, which uses state-action pairs of each agent [23]. As policy, they employ modular, deep attention networks. Agunloye et al. use multi-agent generative adversarial imitation learning to learn swarm behaviors for a swarm of UAVs [24]. In order to generate homogeneous policies, the authors employed MA-PPO with parameter sharing. The collective behavior is modeled through agent-level features that compute the distances to all relevant elements in the environment.

One main challenge in imitation learning for swarm robotics remains the way demonstrations are provided. Most of the above mentioned works assume that a policy already exists that generates the desired behavior. In practical terms, this leads to a bootstrapping paradox: If the policy already exists, there is no need to imitate it and if it cannot be generated easily, the imitation learning process has no access to expert demonstrations. Only few works address directly the problem of a human operator providing demonstrations. Notably, Gharbi et al. [20] and Szpirer et al. [21] use demonstrations that include only the final position of all robots in the environment. These demonstrations can easily be provided by a human operator prior to the imitation learning process. Mattson and Brown tackled the problem of giving demonstrations to a whole swarm by designing a round-robin system [14]. The human operator would control a single robot, while the remaining robots in the swarm were controlled by random policies. The human-controlled robot would then use behavior cloning to derive a policy mimicking the demonstration. In the following steps, the human would switch through the other robots, providing demonstrations for each of them while the remaining robots use their learned policies.

III. METHODOLOGY

In order to apply imitation learning techniques to swarm robotics, a few considerations need to be made. Features can be defined either on the agent-level or on a collective level. While agent-level features are more closely connected to the policy, they cannot capture all properties of the swarm. Therefore, we focus on swarm-level features (see Sec. III-B). This allows us to model the problem as a single-agent imitation learning problem (see Sec. III-D).

TABLE I: The simplified model for the TurtleBot 4 robot.

Sensor	Range	Dim.
Speed $ v $	$[0, 0.31m]$ m/s	1
LiDAR	$[0, 200]$ cm	5
Color	$\{0; 1\}$	3
Bumper	$\{0; 1\}$	1
Actuator	Range	Dim.
Linear vel. v	$[0, 0.31]$ m/s	1
Angular vel. ω	$[-1.9, 1.9]$ rad/s	1

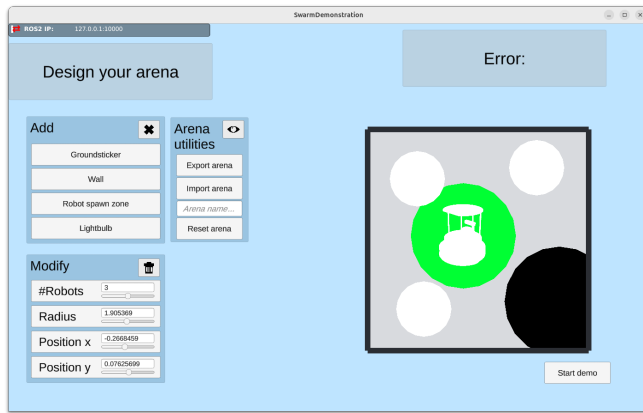
A. Robot model

We provide demonstrations and design policies for a swarm of three TurtleBot 4 robots¹. These robots are commonly used in robotics research and provide several sensors, including infra-red proximity sensors, a 2-d LiDAR, and a bumper for collision detection. In order to reduce the dimensionality of the sensor data, we consider a simplified robot model (see Table I). The robot is modeled as a differential drive agent with linear velocity v and angular velocity ω . For sensing, the robot has access to its own current absolute speed $|v|$. It can also sense neighboring robots through its LiDAR. The LiDAR readings are aggregated into five equal sectors (front-right, right, front-left, left, back) that cover 72° each. For each sector, only the closest robot is considered. If a sector does not contain any other robot, the reading is capped at 200 cm. In order to be able to react to different zones in the environment, we model that the robots are equipped with ground sensors. The ground sensors can detect three distinct gray-scale colors: black, white and gray. The ground color is provided as a one-hot encoded, three-dimensional vector. Lastly, robots are able to access their bumper to detect collisions with other robots or the worlds.

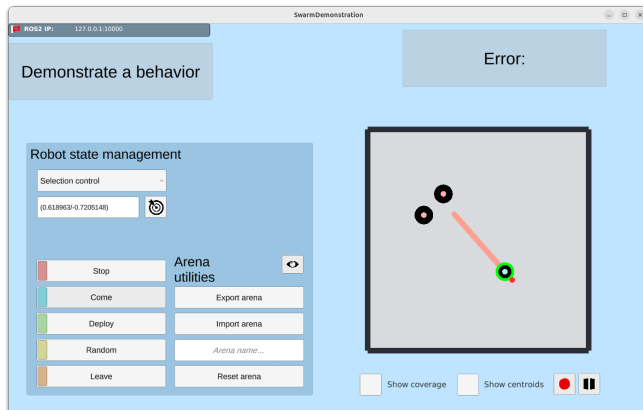
B. Features for collective behavior

We select five features to describe the collective behavior: average speed, grouping, coverage, color travel time, and color visit frequency. These features are computed at every time step during the execution of the mission. *Average speed* is defined as $\frac{1}{N} \sum_{i=1}^N |v_i|$, with N as the swarm size and v_i as the linear velocity of robot i . The *grouping* feature describes the average distance of the robots to the swarm center. It is computed by $\frac{1}{N} \sum_{i=1}^N |x_i - \bar{x}|$ with N as the swarm size, x_i as the current position of robot i and \bar{x} as the center of mass position of the swarm. In order to compute *coverage*, the arena is divided into a 4×4 grid. Each tile counts the time since it was last visited by a robot. The grouping and coverage metric are taken from [25]. We also include two metrics to describe the behavior of the robots with respect to the colored zones in the environment. For the *color visit frequency*, we keep track of the sequence of ground colors that each robot visits. As gray is the default color of the environment, we focus only on white and black floor patches. Every time a robot enters a colored patch, we increment a counter based on the current color and the previous visited ground color. This results in four counters,

¹<https://clearpathrobotics.com/turtlebot-4/>



(a) Arena building stage



(b) Demonstration building stage

Fig. 1: Two screen captures from our demonstration tool. The user can build an experimental arena and provide demonstrations of swarm behaviors.

one for each combination of white-white, white-black, black-white, and black-black. The *color travel time* is defined as the average time it takes a robot to change between floor patches of any color. This is one feature describing the average time of all of the four possible combinations. This results in a total of 23 swarm-level features for every time step.

C. Demonstration Tool/Simulator

In order to provide demonstrations, we develop a demonstration tool (see Fig. 1). The tool is based on the Unity game engine² and is available for download³. The demonstration tool allows to design different environments for the swarm (see Fig. 1a). The user can modify the $4\text{ m} \times 4\text{ m}$ arena by placing walls and black or white ground patches. In addition, the user can specify spawn zones for the robots. In the beginning of the experiment, the robots will randomly placed within these zones. These configured arenas can be exported and imported so experiments can be easily repeated.

After designing the environment for the swarm, the user can provide demonstrations (see Fig. 1b). Using the demonstration tool, the user can give high-level commands to

the swarm or individual robots in it. Prior research has shown that high-level interaction commands result in better demonstrations than when users are presented with low-level control [26]. The available interaction commands are adapted from Kolling et al. [27]. Kolling et al. define two control modes: *selection control* and *beacon control*. In selection control, the user can select which robots shall change their behavior by clicking on them. Beacon control is an environmental control method, where the user does not directly change the behavior of the robots. Rather, they can place beacons, trigger zones that are connected a specific behavior. Whenever a robot enters the beacon, it immediately switches to the connected behavior of the beacon.

Kolling et al. defined seven behaviors that each robot could execute: *Stop*, *Random*, *Come*, *Leave*, *Heading*, *Rendezvous*, and *Deploy*. However, not every behavior was used equally often by the operators. Therefore, we did not implement *Rendezvous* and *Heading*. The remaining behaviors were implemented as follows: When executing the *Stop* behavior, the robot is standing still. In the *Random* behavior, the robot performs a random walk in the form of ballistic motion. The robot moves straight with maximal speed until it encounters an obstacle. It then turns in place for a random number of time steps before continuing to move straight. For the *Come* and *Leave* behaviors, a target location within the arena needs to be chosen. In the *Come* behavior, the robot drives with maximal speed towards the target. When approaching the target, it progressively slows down. When the target is within a radius of 20 cm, the robot stops. The *Leave* behavior repels the robot from the target. On collision it switches its behavior to *Random*. In the *Deploy* behavior, the robot computes a Voronoi diagram with each robot position as site. The robot drives with maximal speed to the *centroid* of its assigned cell. On collision with other robots, it changes its direction slightly to the right, to avoid further collisions.

In order to collect the user demonstrations, the demonstration tool is able to provide a light-weight simulation of the robot behavior. In order to minimize the difference in simulation contexts between demonstration and design, we also use the demonstration tool as simulator for the imitation learning process. Communication between the GAIL library and the demonstration tool is facilitated using ROS 2 and the Unity Robotics Hub⁴.

D. SwarmGAIL

We reduce the robot swarm imitation learning problem to a single-agent generative adversarial imitation learning problem. For our imitation learning algorithm, we use a slightly adjusted version of GAIL [12] (see Fig. 2). The algorithm learns a single policy (generator) that controls each robot in the swarm in a round-robin fashion. The policy takes as input the observations generated by the robot model (see Table I) and outputs the actuator commands for the focal robot. As the policy is decentralized and only depends on

²<https://unity.com/products/unity-engine>

³<https://github.com/CPS-Konstanz/SwarmGAIL>

⁴<https://github.com/Unity-Technologies/Unity-Robotics-Hub>

Concept for GAIL Operating on Swarms

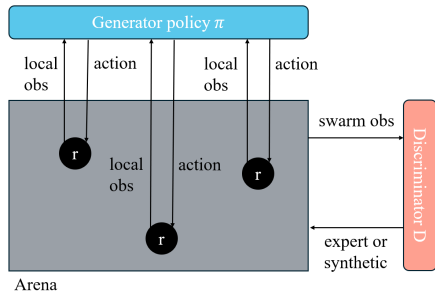


Fig. 2: Information flow in our SwarmGAIL implementation. The discriminator receives swarm-level features as observation. The policy is used in a round-robin fashion to provide actions to all robots based on their local observations.

the observations of a single robot, we can later deploy a copy of the policy on each robot.

The discriminator aims to distinguish between demonstrated and generated behavior. As input, we provide feature-action pairs (with the swarm-level features defined in Section III-B). Notably, the features do not include the local observations of the current focal robot. This ensures that the discriminator only tries to make a classification based on the swarm behavior and not inadvertent differences in the observations between demonstrations and generated behaviors.

We implemented this adjusted GAIL version in Python 3 using the `imitation` library [28]. The discriminator is a multi-layer perceptron with two hidden layers of size 32 each. The policy is a multi-layer perceptron with two hidden layers of size 64 each and is trained using PPO [13]. For technical reasons, the policy and discriminator expect the same observation space. We therefore consider a joint observation vector that combines local observations and swarm features. We overwrite the unnecessary parts of the observation vector with 0 before supplying the vector to either the discriminator or the policy.

IV. EXPERIMENTAL SETUP

A. Missions

To test our imitation learning framework, we consider six missions: STANDING STILL, FULL SPEED, CONTROLLED SPEED, AGGREGATION, DISPERSION and FORAGING (see Fig. 3). For each mission, we define a general strategy to provide demonstrations and, additionally, a reward function that we assume is maximized by our demonstrations.

The missions STANDING STILL, FULL SPEED, CONTROLLED SPEED, AGGREGATION and DISPERSION take place in an empty arena. The robots spawn approximately in the middle of the environment. In STANDING STILL, the robots should not move forward (stationary rotation is allowed). The demonstration is performed by selecting all robots and assigning the *Stop* behavior. The reward function is defined as $R_s(t) = -\bar{v}(t)$, where $\bar{v}(t)$ is the average linear velocity of the swarm at time step t . In FULL SPEED,

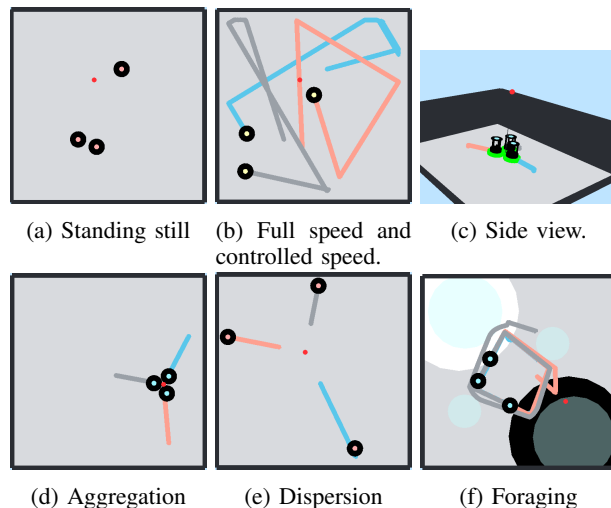


Fig. 3: Representations of the demonstrated behavior for all considered missions. In the case of FULL SPEED and CONTROLLED SPEED, the trajectories taken by the robots look visually similar, but the velocity of the robots differs.

robots should move as fast as possible. Therefore, the reward function for this mission is defined as $R_f(t) = \bar{v}(t)$. For the demonstrations, all robots are selected and we apply the *Random* behavior to them. For CONTROLLED SPEED, the robots should move at a constant velocity $v_{target} = 0.1$ m/s. As none of the behaviors in our demonstration tool allows to explicitly set the linear velocity, we reduced the maximum allowed velocity for the demonstrations. Then we selected all robots and applied the *Random* behavior to them. The reward function for this mission is defined as $R_c(t) = |v_{target} - \bar{v}(t)|$.

In AGGREGATION, the robots should aggregate in the arena, but no particular spot is required for the aggregation. The reward function is defined as $R_a(t) = 1/\bar{g}$, where \bar{g} is the grouping feature, i.e., the average distance of the robots to the center of mass of the swarm. Demonstrations were provided by selecting all robots and assigning them the *Come* behavior with the target being chosen as approximately the center of mass of the swarm. Similarly, in DISPERSION, the robots should spread out as far as possible in the arena. The reward function is defined as $R_d(t) = -1/\bar{g}$. Demonstrations are provided by selecting all robots and assigning them the *Leave* behavior. FORAGING is the most complex behavior that we considered. The environment contains a white area (the nest) and a black area (the source) and the robots spawn in the middle of the arena. Robots are tasked to collect (virtual) items from the nest and deposit them at the source. Retrieval and deposition of the items happen automatically, as the robots enter the respective zones. Every successfully retrieved or deposited item is rewarded with a score of 10. Demonstrations were provided through beacon control. Two beacons were placed in the gray area. When entering these beacons, the robot would switch to the *Come* behavior with the target within the nest and source, respectively. The nest and source also contain beacons that, when entered, would

switch the robots to the *Come* behavior, with the target being the way point to the respective next goal. Initially, the robots are selected and assigned a *Come* behavior towards the closest region of interest (way point, nest, or source).

We are aware, that these are basic missions and the desired collective behavior can be achieved using approaches other than imitation. In fact, implementations of these behaviors exist readily in the literature [29]. We choose to focus on these basic missions to reduce the risk of unsuccessful demonstrations. This allows us to interpret the performance as the result of our proposed framework, rather than of the quality of the acquired demonstrations.

B. Real-world environment

In addition to the simulated missions, we also consider experiments in a real-world environment. The robot swarm operates in a $2\text{ m} \times 4\text{ m}$ arena, which is enclosed by walls. Each robot executes its own decentralized policy. In order to prevent damage to the robots and the environment, we implemented a hardware protection layer. Once a (potential) collision is detected, the hardware protection layer will inhibit the output of the policy and rotate in place until the obstacle is cleared. Furthermore, the robots are tracked through a Vicon motion capture system. This allows us to track the positions and orientations of the robots inside the arena and to compute the rewards obtained in these experiments.

C. Experimental Protocol

For each mission, we conduct five independent experimental runs. In each run, we provide one human-operated demonstration that is used to generate five different rollouts, resulting in five trajectories of feature-action pairs. We provide demonstrations following the general strategies described above, allowing for refinement of the exact demonstrations until we were satisfied with the swarm performance. Each demonstration takes 51 seconds. In our experiments, this proved to be a good tradeoff between time for the desired behavior to arise and total time necessary to evaluate all policies. The set of five demonstration rollouts was then provided as the expert demonstrations to the imitation learning framework. In order to reduce the time necessary to evaluate the policies, we reduced the control frequency of the robots in simulation to 1 Hz. To track the development of the imitation learning process, we log information related to the discriminator, as well as a copy of the current policy in regular intervals (for a total of 100 logged policies per run). To assess the quality of the imitated policy, it is then evaluated five times in simulation. Additionally, we transfer it to the real robots and evaluate it once in the real world. The control frequency for the real robots is set to 10 Hz.

Additionally, we also consider imitation learning from trained policies. To that end, we perform a standard reinforcement learning process for each mission using PPO and the policy architecture described in Sec. III-D. The generated policy is then evaluated five times in simulation to generate five policy rollouts (analogous to the five human-operated

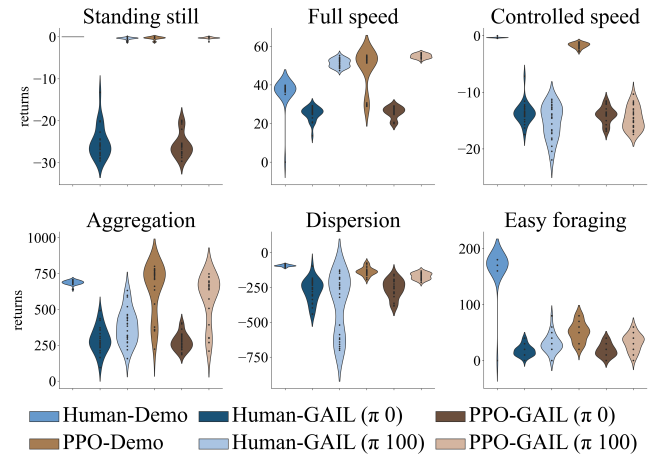


Fig. 4: Violin plots of the return (cumulative reward) of all evaluations across all considered missions. Colors indicate the source of the demonstrations (blue from human-operated, brown from PPO-trained ones). Dark colors represent the initial policies of each imitation learning experiment, light colors the final policy. Black circles correspond to one evaluation of one policy/demonstration. The higher the return, the better the performance.

demonstrations). These rollouts are then used as inputs into an imitation learning process under the same experimental protocol as for human-operated demonstrations. In particular, we also repeat this process five times per mission, leading to five different learned policies, one for each repetition.

All code, demonstrations, policies, performance data and videos are available online⁵.

V. RESULTS

A. Manual demonstrations vs. trained policies

Fig. 4 shows the performance of the rollouts of all demonstrations. As a general trend, using manual control or a trained policy results in similarly good performing behaviors. This is in so far noteworthy, as the human demonstrations do not explicitly aim to maximize the reward function. Differences in the quality are often results from particularities of the mission and demonstrations. For example, in *STANDING STILL*, the manual demonstration easily reaches a score of exactly 0, as it uses the *Stop* behavior on all robots from the beginning. The PPO-trained policy, however, has troubles setting the linear velocity to exactly 0.0 m s^{-1} all the time. As a result, some of the robots move from time to time, accumulating small penalties.

In the mission *FULL SPEED*, contrarily, the PPO trained policies usually discover strategies that are not available to the human demonstrator. Our demonstrations make use of the *Random walk* behavior, which has the robots performing a ballistic motion behavior. The robots move straight with their maximal velocity until they encounter a wall, at which point they turn in place. The turning in place results in short

⁵<https://doi.org/10.5281/zenodo.18877502>

periods of time where the linear velocity of this robot is 0.0 ms^{-1} , therefore not contributing to the reward. In the PPO-trained policies, however, the robots usually drive in a circular motion. This has the effect of keeping the linear velocity (nearly) constant, while the angular velocity is set very high to avoid collisions. Of special note is also the mission FORAGING. The manually provided demonstrations usually achieve scores of more than 150. However, no rollout of the PPO-trained policy results in a score of at least 100. This observation supports our belief that a greater focus should be put on human demonstrations. Especially in more complex missions (of which FORAGING is still a simple representative), human demonstrations can easily eclipse those behaviors generated by other methods.

Lastly, it should be noted that the human demonstrations showed less variance in their performance, compared to the PPO-trained policies. This happens for two reasons: First, it can happen that some of the five PPO-trained policies fail to adequately address the mission. This was, however, never a problem with the manual demonstrations. Second, the performance of the same PPO-policy varies more across repeated rollouts than the performance of the human operator across repeated demonstrations. Only in the mission FORAGING, do the human-operated demonstrations show greater variance than the ones generated from a PPO-trained policy.

B. Generative adversarial imitation learning

Fig. 4 shows also the performance of the generative adversarial imitation learning process. Besides the performance of the demonstrations (human and PPO-trained), we also show the performance of the initial policy (π_0) and the final generated one (π_{100}). As initial policies are randomly generated, their performance only serves as a baseline to highlight the degree to which the performance improved.

For the missions STANDING STILL and FULL SPEED, we observe that initial policies are, as expected, relative poor. The final trained policies, however, match the quality of the demonstrations, and in the case of FULL SPEED even exceed the manual demonstration. The learned policies make use of the fact that they can control both the linear and angular velocity, while manual demonstrations relied on ballistic motion. As previously mentioned, the ballistic motion accrues small penalties for turning in place after collisions. In the imitating policies, robots however set their angular velocity to avoid collisions. This strategy is similar to the one used in PPO-trained policies. Yet, the policies trained from human-operated demonstrations, turn less than the PPO-trained policies, which essentially turn in place. This is due to the fact that the ballistic motion covered the full arena, affecting swarm-level features such as *grouping* or *coverage*. Therefore, also the imitating policies cover the complete arena. Accordingly, the policies imitating the PPO-trained demonstrations learned also to turn nearly on the spot, mimicking closely their demonstrations. This shows that the imitation learning process was able to pick up on these idiosyncrasies in the demonstrated behavior, which are reproduced in the corresponding imitating policies. Yet, the

fact that GAIL was not able to correctly imitate the turning-in place of ballistic motion warrants further investigation, even if on paper it resulted in a better performing behavior.

In the case of CONTROLLED SPEED, the imitating policies do not reach the same performance as the demonstrations. In fact, when compared to the initial policies, the imitation process seems to have affected the performance negatively. Visual inspection of the initial and final policies indicate that both appear to rely on controlling their linear velocity through randomly stopping and moving. This incurs high penalties, as the average velocity is then most of the time either too low or too high. The human-operated and PPO-trained demonstrations, however, move constantly at nearly the target velocity. It is not clear, why the demonstrated behavior did not emerge in either imitation learning process.

In AGGREGATION, the GAIL runs, that used the rollouts from PPO-trained policies as demonstrations, were able to learn to imitate the behavior well. The policies imitating human demonstrations do improve over the random initial ones but do not score as well as the originals. Visual inspection of the generated behaviors reveals that robots are mostly stationary and only try to slowly approach each other. The score of the behavior is therefore highly contingent on the initial starting positions of the robots. This might be caused by the interplay of swarm-level features and manual demonstrations. The human operator provides a quick aggregation, after which the robots remain stationary. This results in most of the features to develop similarly to a stationary swarm (except for the early increased *average speed*). It seems difficult for the imitation learning process to copy this aggregation but it seems to focus mostly on the longer, stationary part of the demonstration. In the PPO-trained demonstrations and the resulting imitating policies, however, the aggregation is not static but dynamic (except for one pair of policies that also learn the static aggregation). This generates a more dynamic representation in the swarm-level features, which apparently is more helpful to the imitation learning process.

In the mission DISPERSION, the imitation learning improves only slightly over the random initial policies. They, however, show less variance in their performance (with the exception of one policy that appears to have learned an aggregation behavior). This can be explained by the fact that the initial policies do not contain very structured motion. As a result, their score is mostly dependent on the initial position. The learned policies provide more structured behavior that allows the robots to separate, should they have started too close together. Yet, as the LiDAR readings of the robots are limited to 2 m, the robots can only separate themselves until they don't perceive each other. The human demonstrator could go one step further and position the robots maximally spread out in the arena.

In the last considered mission, FORAGING, the human demonstrations clearly outperform the imitating policies. Our initial belief was, that the robots might learn to move in circles/spirals that overlap with the source and the nest (and the manual demonstrations draw inspiration from this

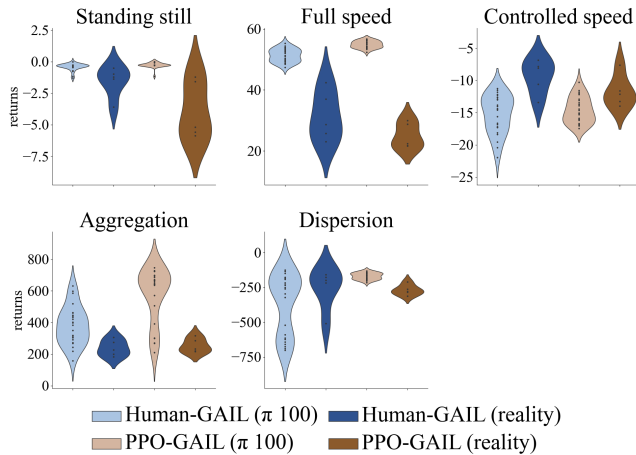


Fig. 5: Violin plots of the returns (cumulative reward) of imitated policies in simulation (light colored) and reality (dark colored). Colors indicate the source of the demonstrations (blue from human-operated, brown from PPO-trained ones). Black circles correspond to one evaluation of one policy/demonstration. The higher the return the better the performance.

strategy as well with their way points). However, the initial policies usually generate a slow, randomly moving behavior. The final trained policies (both from human-operated demonstrations and PPO-trained ones) exhibit a much faster, but nonetheless randomly moving behavior. The increased speed of movement makes it more likely to retrieve and store items, but those occurrences remain accidental. As also the PPO-trained demonstrations fail to exhibit any behavior that appears to be meaningfully different from a random walk, we believe this might be an indication that the environment lacks structure that the robots can exploit for navigation. Additionally, our swarm-level features related to the colors might have further complicated the imitation learning process. The manual demonstrations easily manage to navigate between floor patches, resulting high *visit frequencies* and low *color travel times*. As navigation for the learned policies is much harder, they are unable to match these features. As a result, the discriminator can very easily distinguish between demonstrated and learned behaviors. Therefore, the reward landscape for the policy improvement is flat, as no possible improvement will overcome the discriminator.

C. Real-world validation

Fig. 5 shows the performance of the experiments with real robots. We excluded the mission FORAGING from consideration for two reasons. First, the real TurtleBot 4 robots did not have a ground sensor that would have been necessary to perform this mission. While we could have provided a virtual ground sensor, we also changed the arena size. The policies were not particularly well-performing in simulation and the added change of arena size would probably have impacted the performance even further.

It can be seen that the policies generated from manual

and PPO-trained demonstrations perform similarly, when evaluated on the real robots. This is not surprising, as these policies already performed relatively similar in simulation. Notably, visual inspection of the robot swarm reveals that each mission results in clearly distinguishable and recognizable behavior (with the exception of the one dispersion policy that resulted in an aggregation-like behavior). This further proves that our framework was able to successfully imitate the demonstrated behaviors and the generated behaviors remain qualitatively robust to the reality gap.

Quantitatively, the policies are affected in different ways by the reality gap. The biggest impact on behavior is generated by the hardware protection layer. In some missions, such as AGGREGATION and FULL SPEED, the hardware protection layer prevents collisions that are considered an acceptable part of the behavior in simulation. As a result robots slow down and turn away or stop farther from each other than would happen in simulation. Interestingly, this same effect seems to have been beneficial for the mission CONTROLLED SPEED, where robots performed better in reality than in simulation. The performance in STANDING STILL and DISPERSION was not significantly affected by the hardware protection layer, as robots did not encounter each other often during the experimental runs. This leads us to believe that our method might transfer well to real robots, if the hardware protection layer would have been appropriately modelled in the simulation.

VI. CONCLUSION

In this work, we have presented an imitation learning framework for robot swarms. We model the imitation learning problem as a single-agent problem and solve it using GAIL. The discriminator classifies samples according to five swarm-level features and the policy controls robots according to their own local observations. We have also provided a tool that allows us to provide human-operated demonstrations. Our results show that for all missions except CONTROLLED SPEED and FORAGING, the imitation learning process was able to learn qualitatively meaningful behaviors that performed similarly well as the provided demonstrations. In addition, we have tested the policies on real robots. The transfer of the policies went relatively well with the behaviors being still visually recognizable on the real robots. However, the hardware protection layer on the real robots was not modeled in simulation and therefore impacted the performance of the transferred policies.

Our work has shown, however, also some limitations of our approach. Despite considering relatively simple missions, in some cases, the demonstrated behavior could not always be accurately imitated. We believe that this is partially due to our ad-hoc selected swarm-level features, which impacted the imitation learning process, however, more research on identifying the limiting factors is needed. Another potential factor is the increased performance variance by both human and PPO-trained demonstrations in the more complex missions, such as FORAGING. We believe that more work is needed in order to provide reliable demonstrations for these complex

missions. In addition, future work should be devoted to better understanding how the features impact the imitation learning process (e.g., by performing ablation studies) and how researchers can systematically select these features to ensure proper performance of the imitated behaviors. The fact that we used the same features as input for the discriminator and the reward function is another decision that warrants further consideration. This could lead to instances of information leaking, in which the discriminator implicitly learns which features were associated with the rewards function assumed to be maximized by the demonstration. In future work, a more decoupled approach to feature-space and reward function should be considered. This applies in particular, when moving to more complex missions.

Despite these limitations, our approach also opens up new avenues for research. We believe that it would be of interest to compare this approach against other methods from the literature, to source demonstrations from experts with varying levels of expertise in swarm robotics, and to study how well the trained policy scales to different swarm sizes.

ACKNOWLEDGMENT

We thank Mahabaleshwar Ammu for their help in setting up the experiments with real robots.

REFERENCES

- [1] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: a review from the swarm engineering perspective," *Swarm Intell.*, vol. 7, no. 1, pp. 1–41, 2013.
- [2] M. Dorigo, M. Birattari, and M. Brambilla, "Swarm robotics," *Scholarpedia*, vol. 9, no. 1, p. 1463, 2014.
- [3] M. Schranz, M. Umlauf, M. Sende, and W. Elmenreich, "Swarm robotic behaviors and current applications," *Front. Robot. AI*, vol. 7, p. 36, 2020.
- [4] M. Dorigo, G. Theraulaz, and V. Trianni, "Swarm robotics: past, present, and future [point of view]," *Proceedings of the IEEE*, vol. 109, no. 7, pp. 1152–1165, 2021.
- [5] G. Francesca and M. Birattari, "Automatic design of robot swarms: achievements and challenges," *Front. Robot. AI*, vol. 3, no. 29, pp. 1–9, 2016.
- [6] J. Kuckling, "Recent trends in robot learning and evolution for swarm robotics," *Front. Robot. AI*, vol. 10, 2023.
- [7] M. Birattari, A. Ligot, D. Bozhinoski, M. Brambilla, G. Francesca, L. Garattoni, D. Garzón Ramos, K. Hasselmann, M. Kegeleirs, J. Kuckling, F. Pagnozzi, A. Roli, M. Salman, and T. Stützel, "Automatic off-line design of robot swarms: a manifesto," *Front. Robot. AI*, vol. 6, p. 59, 2019.
- [8] J. Lehman and K. O. Stanley, "Abandoning objectives: evolution through the search for novelty alone," *Evol. Comput.*, vol. 19, no. 2, pp. 189–223, 2011.
- [9] S. Doncieux, N. Bredeche, J.-B. Mouret, and A. Eiben, "Evolutionary robotics: what, why, and where to," *Front. Robot. AI*, vol. 2, p. 4, 2015.
- [10] M. Divband Soorati and H. Hamann, "The effect of fitness function design on performance in evolutionary robotics: the influence of a priori knowledge," in *GECCO 2015*, S. Silva, Ed. New York, NY, USA: ACM, 2015, pp. 153–160.
- [11] A. Y. Ng, D. Harada, and S. J. Russell, "Policy invariance under reward transformations: Theory and application to reward shaping," in *Proceedings of the Sixteenth International Conference on Machine Learning*, ser. ICML '99. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999, p. 278–287.
- [12] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, 2016.
- [13] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [14] C. Mattson and D. S. Brown, "R2BC: Multi-Agent Imitation Learning from Single-Agent Demonstrations," 2025. [Online]. Available: <https://hitl-robot-learning.github.io/pdfs/r2bc.pdf>
- [15] I. M. Rahman, S. White, K. Crockett, Y. Gu, D. A. A. Dutra, and G. A. S. Pereira, "Imitating Swarm Behaviors by Learning Agent-Level Controllers," in *2023 American Control Conference (ACC)*. San Diego, CA, USA: IEEE, May 2023, pp. 7–13. [Online]. Available: <https://ieeexplore.ieee.org/document/10156561/>
- [16] S. Zhou, M. Phielipp, J. A. Sefair, S. I. Walker, and H. B. Amor, "Clone Swarms: Learning to Predict and Control Multi-Robot Systems by Imitation," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov. 2019, pp. 4092–4099, arXiv:1912.02811 [cs]. [Online]. Available: <http://arxiv.org/abs/1912.02811>
- [17] K. Alharthi, Z. S. Abdallah, and S. Hauer, "Automatic extraction of understandable controllers from video observations of swarm behaviors," in *International Conference on Swarm Intelligence*. Springer, 2022, pp. 41–53.
- [18] —, "Ghost swarms: Learning swarm rules from environmental changes alone," in *European Conference on Genetic Programming (Part of EvoStar)*. Springer, 2025, pp. 1–17.
- [19] A. Šošić, W. R. KhudaBukhsh, A. M. Zoubir, and H. Koepl, "Inverse Reinforcement Learning in Swarm Systems," in *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*, 2017. [Online]. Available: <https://ifaamas.org/Proceedings/aamas2017/pdfs/p1413.pdf>
- [20] I. Gharbi, J. Kuckling, D. Garzón Ramos, and M. Birattari, "Show me what you want: Inverse reinforcement learning to automatically design robot swarms by demonstration," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. Piscataway, NJ, USA: IEEE, 2023, pp. 5064–5070.
- [21] J. Szpírer, D. G. Ramos, and M. Birattari, "Automatic design of robot swarms that perform composite missions: an approach based on inverse reinforcement learning," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 5791–5798.
- [22] W. Li, M. Gauci, and R. Groß, "Turing learning: a metric-free approach to inferring behavior and its application to swarms," *Swarm Intelligence*, vol. 10, no. 3, pp. 211–243, Sep. 2016. [Online]. Available: <http://link.springer.com/10.1007/s11721-016-0126-1>
- [23] Y. Wu, T. Wang, T. Liu, Z. Zheng, D. Xu, and X. Peng, "Adversarial imitation learning with deep attention network for swarm systems," *Complex & Intelligent Systems*, vol. 11, no. 1, p. 26, Jan. 2025. [Online]. Available: <https://link.springer.com/10.1007/s40747-024-01662-2>
- [24] A. O. Agunloye, S. D. Ramchurn, and M. D. Soorati, "Learning to Imitate Spatial Organization in Multi-robot Systems," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Abu Dhabi, United Arab Emirates: IEEE, Oct. 2024, pp. 3463–3469. [Online]. Available: <https://ieeexplore.ieee.org/document/10801662/>
- [25] L. G. Jeston-Fenton, S. Abpeikar, and K. Kasmarik, "Visualisation of Swarm Metrics on a Handheld Device for Human-Swarm Interaction," in *2022 26th International Conference Information Visualisation (IV)*. Vienna, Austria: IEEE, Jul. 2022, pp. 149–154. [Online]. Available: <https://ieeexplore.ieee.org/document/10017712/>
- [26] A. Hussein, H. Nguyen, and H. A. Abbass, "Machine teaching in swarm metaverse under different levels of autonomy," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 383, no. 2289, p. 20240149, 2025.
- [27] A. Kolling, S. Nunnally, and M. Lewis, "Towards human control of robot swarms," in *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Mar. 2012, pp. 89–96, iSSN: 2167-2148. [Online]. Available: <https://ieeexplore.ieee.org/document/6249619>
- [28] A. Gleave, M. Taufeeque, J. Rocamonde, E. Jenner, S. H. Wang, S. Toyer, M. Ernestus, N. Belrose, S. Emmons, and S. Russell, "imitation: Clean imitation learning implementations," arXiv:2211.11972v1 [cs.LG], 2022. [Online]. Available: <https://arxiv.org/abs/2211.11972>
- [29] T. K. Kaiser, M. J. Begemann, T. Plattenteich, L. Schilling, G. Schildbach, and H. Hamann, "ROS2SWARM – a ROS 2 package for swarm robot behaviors," in *ICRA 2022*. IEEE, 2022, pp. 6875–6881.