

Satellite Navigation and Coordination with Limited Information Sharing

Sydney Dolan*

Siddharth Nayak*

Hamsa Balakrishnan

Massachusetts Institute of Technology

SYDNEYD@MIT.EDU

SIDNAYAK@MIT.EDU

HAMSA@MIT.EDU

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

We explore space traffic management as an application of collision-free navigation in multi-agent systems where vehicles have limited observation and communication ranges. We investigate the effectiveness of transferring a collision avoidance multi-agent reinforcement (MARL) model trained on a ground environment to a space one. We demonstrate that the transfer learning model outperforms a model that is trained directly on the space environment. Furthermore, we find that our approach works well even when we consider the perturbations to satellite dynamics caused by the Earth’s oblateness. Finally, we show how our methods can be used to evaluate the benefits of information-sharing between satellite operators in order to improve coordination.

Keywords: transfer learning, multi-agent reinforcement learning, graph neural networks, space traffic management

1. Introduction

There are an estimated 8,800 satellites and over a million pieces of debris in the sky today ([Space Debris Office, 2022](#)); by 2030, there will be an estimated 150,000 active satellites in space ([O’Callaghan, 2022](#)). The sheer numbers of objects in orbit and the resulting potential collisions will likely make current approaches untenable, and make autonomous decision-making an essential characteristic of space traffic management (STM) in the future ([Hobbs et al., 2020](#)).

Multi-agent reinforcement learning (MARL) has yielded promising results in several settings ([Vinyals et al., 2019](#); [Berner et al., 2019](#)), including for navigation and collision avoidance problems ([Lowe et al., 2017a](#); [Yu et al., 2022](#)). Transfer learning has achieved extensive success by leveraging prior knowledge of past learned policies of relevant tasks ([Yang et al., 2020](#); [Khan et al., 2019](#); [Muandet et al., 2017](#)). Inspired by this, we investigate the effectiveness of transferring a ground-based collision avoidance MARL model to space traffic management, or more specifically, for the collision-free navigation of satellites in orbit. We leverage our recent work on a graph neural network (GNN) based architecture for MARL, called InforMARL ([Nayak et al., 2023](#)). We demonstrate that transfer learning from the ground to the space environment is remarkably effective: it achieves better sample complexity and slightly higher rewards than when directly training a model on the space environment. This is despite the two environments being quite different in terms of the underlying dynamics that govern them. We then consider a more refined abstraction of the space environment

* Equal contribution.

that accounts for perturbations in gravitational disturbances due to the Earth’s oblateness and find that transfer learning is still effective.

Finally, we study the role that information-sharing plays in satellite operator decision-making. Operators are hesitant to share information about their satellites for a number of reasons, including proprietary and security concerns (Rendleman and Mountin, 2015). While third-party screening services may scan for potential collisions between satellites belonging to different operators, these services do not have access to the high-quality state information known to the operator of a satellite. Consequently, these screening services often have a high false alarm rate, with a detrimental impact on the trust placed on their alerts by satellite operators (Hiles et al., 2021). Miscommunications between operators has resulted in numerous near-misses between satellites in orbit (Foust, 2021; Grupen et al., 2021). Motivated by these observations, we use our model to assess the value of sharing orbits and maneuver information among satellite operators.

2. Related work

Collision avoidance of spacecraft has traditionally relied on optimal control approaches. For example, Bombardelli and Hernando-Ayuso (2015) relied on the linear relation between the applied thrust and an object’s relative motion in the collision plane to try to maximize the collision miss distance. Building upon Bombardelli (2013)’s analytical formulation, Salemme et al. (2020) designed low-thrust propulsion collision avoidance maneuvers using a finite-burn arc through an indirect optimal control model. De Vittori et al. (2022) focused on computational efficiency and used an indirect trajectory optimization technique for on-board low-thrust collision avoidance maneuver design. The overarching assumption in these aforementioned works is that there is sufficient information about objects in the environment in order to efficiently maneuver and avoid collisions. By contrast, our work studies situations in which complete information may not be readily available.

In practice, information sharing is very limited in the space domain. The Federal Communications Commission and Department of Commerce solicited feedback from operators about space traffic management data sharing (Rendleman and Mountin, 2015; Weeden et al., 2019); in response, commercial operators expressed a desire to limit the exchange of proprietary information that could give their competitors insight into the capabilities, health, and life of their satellites. It was also suggested that some operators may not have high-quality data to share. The above observations motivate the development of methods that can accommodate scenarios with limited information.

Recent work has used reinforcement learning for spacecraft trajectory optimization, guidance, and control (Izzo et al., 2019; Sullivan and Bosanac, 2020; Cheng et al., 2019; Hovell and Ulrich, 2020). In contrast to these works, we focus on multi-agent coordination among satellites (or their operators) to improve the safety and efficiency of space traffic operations. While there has been much recent work on MARL, including for navigation and collision avoidance problems (Lowe et al. (2017a); Yu et al. (2022); see Nayak et al. (2023) for more references), it has seen limited use in space domain applications (Siew et al., 2022). To the best of our knowledge, our work represents one of the first efforts to use MARL for satellite navigation and collision avoidance.

3. Methods

Figure 1 provides an overview of InforMARL (Nayak et al., 2023), the approach that we adapt for space traffic management applications. We make two main modifications to the original InforMARL

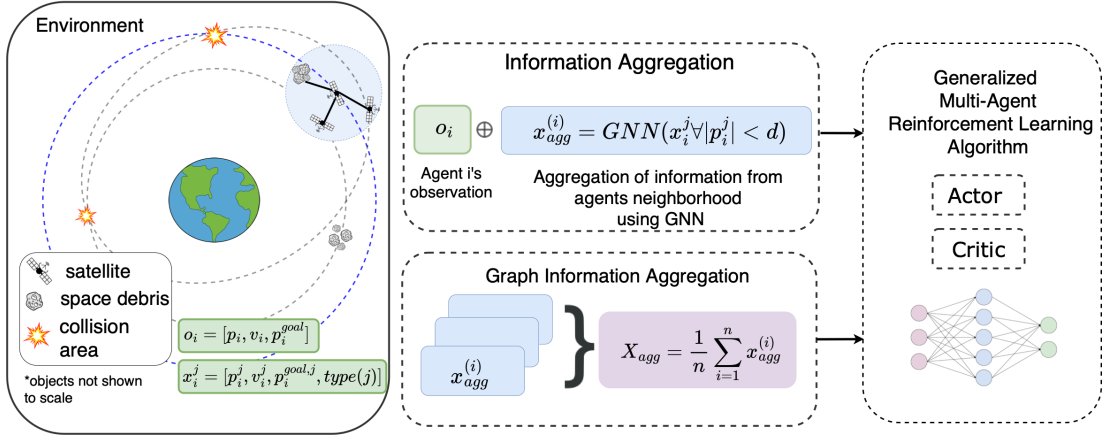


Figure 1: Schematic depiction of InforMARL, the approach that we adapt to the space environment in this paper. The meaning of each icon is shown in the legend on the left of the environment box. Each agent or satellite i receives observations about other satellites and space debris within its sensing radius, denoted by the blue shaded circle. This information is passed into a GNN, which learns compact representations of the relationships between objects. The output of the GNN is denoted $x_{agg}^{(i)}$. (ii) Information aggregation: Each agent’s state and observations (o_i) from a neighborhood d are aggregated to obtain $X_{agg}^{(i)}$. (iii) Graph information aggregation: The vectors of all the agents are averaged to get X_{agg} . (iv) The concatenated vector $[o_i, X_{agg}^{(i)}]$ is input to an RL algorithm to obtain a recommended action. Figure adapted from [Nayak et al. \(2023\)](#).

method. First, we use a new space environment, described in Section 3.1, that reflects the relative dynamics of satellites in orbit. Second, we modify the goal-sharing and goal-setting properties in the information aggregation and graph information aggregation modules of Figure 1, as explained in Section 3.2.

3.1. Environments

In this work, we consider two different environments: (1) A ground environment in which the agents’ dynamics are governed by a double integrator physics model ([Rao and Bernstein, 2001](#)), and (2) a space environment in which the relative motion between two agents (satellites) follows the Clohessy-Wiltshire equations ([Vallado and McClain, 2007](#)). We also consider a modified space environment that accounts for perturbations in the dynamics caused by the oblateness of the Earth.

3.1.1. GROUND ENVIRONMENT: DOUBLE-INTEGRATOR MODEL

For the ground environment in our transfer learning experiments, we rely on a double integrator physics model to simulate the motion of agents “on the ground” ([Rao and Bernstein, 2001](#)). Our environment is a modification of the one used in [Lowe et al. \(2017b\)](#). It corresponds to a 2D space in which agents move based on the following dynamics:

$$m\ddot{x} = -\gamma\dot{x} + f_x \quad (1)$$

$$m\ddot{y} = -\gamma\dot{y} + f_y \quad (2)$$

Here, $f_x(t)$ and $f_y(t)$ represent the x and y components of the total force on an agent at time t . The total force is a sum of the control action from the RL algorithm (see point (iv) in Figure 1) and any collision forces experienced by the agent. m is the mass of each entity, and γ is a damping coefficient. In our ground environment simulations, each agent starts from a random, stationary position. The environment size is a hyperparameter, defining the area available for an agent to move in.

3.1.2. SPACE ENVIRONMENT: CLOHESSY-WILTSHIRE EQUATIONS




Figure 2: The Radial In-track Cross-track (RIC) coordinate frame, centered on the target satellite. The radial direction is defined along the satellite's position vector, and the cross-track direction along its angular momentum vector. The in-track direction completes the triad.

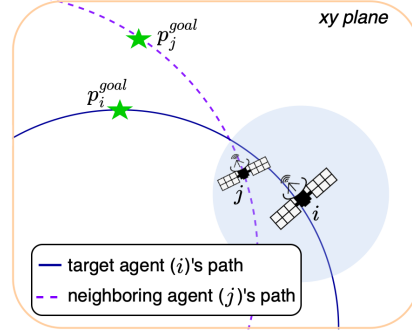


Figure 3: The xy plane highlighted in Figure 2. The shaded circle represents the sensing region of satellite i . Each satellite has its own goal, denoted by p_i^{goal} and p_j^{goal} .

For two objects in orbit, we consider their *relative* motion to be governed by Clohessy-Wiltshire Equations (Vallado and McClain, 2007), as follows:

$$\ddot{x} = 3\omega_n^2 x + 2\omega_n \dot{y} \quad (3)$$

$$\ddot{y} = -2\omega_n \dot{x} \quad (4)$$

$$\ddot{z} = -\omega_n^2 z \quad (5)$$

The above equations consider a localized coordinate system centered around one of the satellites, referred to as the *target*). The target satellite is assumed to have a circular orbit with an orbital rate of ω_n . The coordinates are defined such that x is measured radially outward from the target, y is along the orbit track of the target body, and z is along the angular momentum, as shown in Figure 2.

In general, spacecraft and debris have elliptical (including possibly circular) orbits. For example, suppose the target satellite in Figure 3 has an orbit depicted by the solid line. The orbit of a nearby satellite on the same plane is shown by the dashed line. Suppose the target satellite maneuvers to shift from its original orbit to a new transfer orbit, depicted by the dotted line. Since the coordinate frame is defined relative to the original frame, we can still track the relative motion of both satellites.

For the perturbed model, the dynamic equations are modified to include gravitational disturbances (or perturbations) due to Earth's oblateness:

$$\ddot{x} = (5c^2 - 2)\omega_n^2 x + 2\omega_n c \dot{y} \quad (6)$$

$$\ddot{y} = -2\omega_n c x \quad (7)$$

$$\ddot{z} = (2 - 3c^2)\omega_n^2 z \quad (8)$$

In the above, c is a parameter that reflects the change in orbital rate experienced by the satellite due to these perturbations. The value of c is given by:

$$c = \sqrt{1 + s}, \text{ where } s = \frac{3J_2R_e^2}{8r^2}(1 + 3\cos(2\phi)). \quad (9)$$

In Equation 9, J_2 is a coefficient representing the gravitational effect of a body’s oblateness, R_e represents the Earth’s radius, r represents the radius of the path of the target satellite, and the inclination, ϕ , is the angle of the orbit relative to Earth’s equator. An angle of $\phi = 0^\circ$ represents an equatorial orbit, whereas $\phi = 90^\circ$ represents a polar orbit. We note that substituting $c = 1$ (i.e., $\phi = 54^\circ$) in Equations (6)-(8) results in the Clohessy-Wiltshire equations ((3)-(5)).

A full derivation of the perturbed equations can be found in Roberts and Roberts (2004). While there exist additional environmental perturbations, such as solar radiative pressure or three-body effects, their impacts are magnitudes smaller (Walter, 2018). Consequently, these other perturbations only affect satellite dynamics over time periods that are much longer than the episode lengths considered in this paper, which are of the order of a few hours.

This work focuses on in-plane maneuvers, which means all satellites and debris involved are assumed to lie on the same xy orbital plane, as shown in Figure 2. Similar to the ground environment, the maneuvers are determined by the control action (a force in the xy plane) recommended by the RL algorithm. In both the original Clohessy-Wiltshire equations and the perturbed model, z -dynamics is decoupled from those in the x - and y -directions. While not a topic of investigation in this paper, we believe that, in principle, the same techniques could be applied to the cross-track dimension.

3.1.3. COMPARISON OF GROUND AND SPACE ENVIRONMENTS

We explore the potential of transfer learning for this application because the space environment has complicated dynamics that are much more challenging for the algorithm to learn. From equations 3 and 4, we see that the x and y variables in the Clohessy-Wiltshire equations are coupled. By contrast, the ground equations are independent of one another.

In both cases, we adopt a simplistic reward function similar to the one used in multi-agent particle environment (Lowe et al., 2017b). We assume that at time t , each agent i gets a reward: $r_t^{(i)} = r_{\text{dist},t}^{(i)} + r_{\text{coll},t}^{(i)} + r_{\text{goal},t}^{(i)}$, where $r_{\text{dist},t}^{(i)}$ is the negative of the Euclidean distance to the goal, $r_{\text{coll},t}^{(i)} = -5$ if it collides with any other entity and zero otherwise, and $r_{\text{goal},t}^{(i)} = +5$ if the agent has reached the goal and zero otherwise. The joint reward function is defined as the sum of the individual agent rewards, which encourages cooperation among all agents. It is worth noting that reward function can be further refined, especially in terms of how collisions are penalized. Future research will include imposing larger penalties for collisions, as well as the use of control barrier functions for providing safety guarantees (Cheng et al., 2020).

It should be noted that the scale of distances in the ground environment is in meters (and speeds in m/s), whereas the scale of the space environment is in kilometers (and km/s, respectively). Despite the orders-of-magnitude differences in the units, the numerical values are quite similar in the ground and space environments. Although the control action is the same (1 N) for both environments, the dynamics in the space environment are much more sensitive to external control actions (Hinkel et al.). We investigated additional scaling mechanisms between the two environments, but found that the sample complexity of InforMARL was sufficient to achieve successful transfer learning.

3.2. Graph representation

We use an agent-entity graph as used by Agarwal et al. (2019); Nayak et al. (2023) to represent the interactions between different satellites and debris in the environment. For each agent i , we define an agent-entity graph at each time-step t . This graph is made up of nodes and edges, $g_t^{(i)} \in \mathcal{G} : (\mathcal{V}, \mathcal{E})$, where each node $v \in \mathcal{V}$ is an entity in the environment. Edges, $e \in \mathcal{E}$, are defined when an agent and an entity are within a sensing radius d of one another. For agent-agent interactions, the edges are bidirectional, meaning that the messages are passed back and forth. For agent-non-agent interactions, the edges are unidirectional, meaning that messages are only passed from the non-agent entity to the agent. Translating this to the space environment, a bi-directional edge means that there is open communication between both satellites, while a unidirectional edge means that the satellite has received information about the location of a nearby piece of debris. Each agent i 's local observation $o^{(i)}$ consists of its position and velocity in a global frame of reference and the relative position of the agent's goal with respect to its position.

As depicted in the environment block in Figure 1, each node j in the graph $g^{(i)}$ has node features $x_j = [p_i^j, v_i^j, p_i^{\text{goal},j}, \text{entity_type}(j)]$ where $p_i^j, v_i^j, p_i^{\text{goal},j}$ are the *relative* position, velocity, and position of the goal of the entity at node j with respect to agent i , respectively. If node j is an obstacle or a goal, then it is set to be equivalent to the position $p_i^{\text{goal},j} \equiv p_i^j$. The variable $\text{entity_type}(j) \in \{\text{agent}, \text{obstacle}, \text{goal}\}$ determines the type of entity at node j . In this context, the *goal* entity type is used to indicate each satellite's intentions to maneuver.

To determine the value of information sharing, we rely on two graph variants. The first graph variant is for the case when satellites share their goals. For this graph type, the agents share both their goal $p_i^{\text{goal},j}$ and their state information p_i^j, v_i^j . The second graph variant is for the case when satellites are hiding their goals. For this graph type, each agent node j on the graph is transformed into: $x_j = [p_i^j, v_i^j, \text{entity_type}(j)]$.

4. Numerical experiments

We first demonstrate the scalability of the approach using the space environment for both training and testing, while varying the number of agents. Next, we consider the effectiveness of transferring a model trained in the ground environment to the space one, both in unperturbed and perturbed cases. We then use our model to evaluate the benefits of satellite operators sharing goal information for the purposes of space traffic management.

Agents start at random locations at the beginning of each episode; the corresponding goals are also randomly distributed. Static obstacles are placed randomly in the environment in each episode. The environment is a $2 \text{ km} \times 2 \text{ km}$ area, and the episode length is approximately one hour. Each scenario is initialized with three pieces of debris, which remain at the same relative locations throughout the simulation. The episode continues even after a collision. However, a collision force is exerted on any colliding agent that affects its subsequent dynamics. The overarching objective is for every agent to reach its corresponding goal without colliding with any other entity. Since the episode continues even after a collision, it is possible for there to be multiple collisions in the same episode.

We calculate the following metrics: 1) The total rewards obtained by the agents during an episode using the reward function defined in Section 3.1.3. A higher value corresponds to better performance. 2) The fraction of an episode taken on average by agents to get to their goals, denoted T (lower is better). T is set to 1 if an agent does not reach its goal. 3) Percent of episodes in which all agents get to their goals, denoted $S\%$ (higher is better). 4) The total number of collisions those agents had

in an episode, denoted # col. The lower this metric, the better the performance of the algorithm. Although having fewer collisions is better, some learned policies do not significantly move the agents from their initial orbit, and hence do not get to the goal but have a lower collision rate. In short, the collision rate and the success rate should be considered in conjunction with each other.

It is important to note that the size of the environment considered (4 sq km) along with the number of objects in it (between 6-13) result in unrealistically dense scenarios. The purpose of these experiments is to evaluate the scalability of the methods and the general trends in performance, and not to determine values (e.g., of the collision rates) that are representative of real-world operations.

4.1. Scalability

		Test		
		<i>m</i> = 3	<i>m</i> = 5	<i>m</i> = 10
<i>n</i> = 3	Reward/ <i>m</i>	61.57	60.21	57.78
	<i>T</i>	0.44	0.44	0.43
	(# col)/ <i>m</i>	0.36	0.77	1.41
	<i>S</i> %	98	94	96
<i>n</i> = 5	Reward/ <i>m</i>	60.52	60.52	57.07
	<i>T</i>	0.44	0.44	0.44
	(# col)/ <i>m</i>	0.78	1.28	1.41
	<i>S</i> %	98	98	91

Table 1: Performance metrics obtained by training InforMARL on a space environment with n satellites and testing it on one with m satellites: (a) Total reward obtained in an episode per agent, Reward/ m . (b) Fraction of episode taken on average by agents to reach their goal, T (lower is better). (c) Average number of collisions per agent in an episode, #col/ m (lower is better). (d) Success rate, $S\%$: percentage of episodes in which all agents are able to get to their goals (higher is better).

The first experiment, shown in Table 1, demonstrates the scalability of our algorithm when trained on n agents and tested on m agents. The number of obstacles is held constant to be 3 throughout both all training and evaluation. When evaluating the connectivity throughout the experiments, we found that each agent maintained a connection with at least one entity in the environment more than 90% of the time, meaning that the vast majority of the experiment involves information-sharing between entities. We find that in all scenarios considered, our approach can control the satellites to reach their goals within approximately 44% of the episode length. As expected, the number of collisions per agent increases when there are more satellites (i.e., the environment has become more dense). As mentioned in Section 3.1.3, these values can be improved further by modifying the reward function or by using control barrier functions. A key finding is that the reward per agent remains approximately the same even when the model is trained with $n < m$. Our approach also has robust sample complexity, with a high success rate for unseen scenarios. As the number of objects in space is expected to dramatically increase in the coming years, the scalability and sample efficiency of our technique make it a promising approach to space traffic management.

4.2. Transfer learning

Figure 4 compares the effectiveness of transfer learning (denoted by InforMARL w/ Transfer Learning) relative to other MARL baseline algorithms. The transfer learning model was initialized using




Figure 4: InforMARL with transfer learning learns a similar reward to RMAPPO despite having access to less information. The transfer learning model also outperforms the baseline InforMARL model that was trained only on a space environment.

weights trained on a ground environment, and then trained in the space environment. Consequently, in Figure 4, we offset the plot of InforMARL with transfer learning by 200,000 steps (the number of steps used to train the ground-based model). We trained each algorithm 5 times on 5 separate seeds. The shaded area envelops one standard deviation of the runs.

Three of the baseline methods (RMAPPO (Yu et al., 2022), VDN (Sunehag et al., 2017) and RMADDPG (Lowe et al., 2017a)) use global information, i.e., every satellite in the environment would be required to share their information publicly. We also explored the performance of QMIX (Rashid et al., 2018) and MATD3 (Ackermann et al., 2019), but the performance was similarly poor (see the appendix of Nayak et al. (2023) for results). While the assumption of global information sharing can help determine a performance bound, it is not realistic in practice. By contrast, InforMARL (both with and without transfer learning) uses only local information. Over the training period, InforMARL with transfer learning reaches a similar reward to RMAPPO (Yu et al., 2022), despite needing significantly less information. This finding indicates that the quality of the information, rather than its quantity, is an important driver of performance. Furthermore, we see in Figure 4 that InforMARL with transfer learning outperforms the InforMARL model that was trained from scratch on the satellite environment. These results indicate that ground-based models can be used to accelerate training for space traffic applications. Initial investigations suggest that the poor performance on some instances without transfer learning are because of allocation to weaker computational nodes. InforMARL with transfer learning appears to be less impacted by the specific computational nodes assigned for training. A possible reason is that the initialization with the ground model helps avoid some of the less efficient learning paths, resulting in more consistent training performance. We plan to study this further in future work.

Based on the above results, we believe that the benefits of using InforMARL with transfer learning for such space applications are mainly two-fold: (1) The simplicity of the ground dynamics and the availability of more established ground simulation environments make transfer learning from

ground to space an attractive approach; and (2) InforMARL with transfer learning appears to be less susceptible to the performance of the specific computational nodes that are assigned for training.

4.3. Sensitivity to perturbations

Table 2 demonstrates the performance of InforMARL with transfer learning in the perturbed satellite environment for 3 agents. Once again, the transfer learning model was initialized using weights trained on a ground environment, and then trained in the space environment. We find that even in perturbed environments, the sample complexity of InforMARL allows the algorithm to learn effectively.

In the perturbed model with dynamic equations given by (6)-(8), an inclination of 54° corresponds to the unperturbed Clohessy-Wiltshire equations. When comparing the additional inclination values on the table to 54° , the transfer learning model still performed well, at a comparable level to the model without perturbations (Table 2).

Inclination (ϕ)	0°	28°	45°	54°	63°	72°	81°	90°
Average reward	170.10	166.33	167.07	169.73	168.14	171.78	174.29	169.76
Standard deviation	± 6.32	± 6.18	± 6.09	± 3.07	± 8.85	± 5.45	± 3.02	± 6.10

Table 2: Sensitivity of InforMARL with transfer learning to different perturbations (orbit inclinations), for a three-agent system. The average total reward achieved at the end of training and its standard deviation are presented, based on 5 runs with 5 different random seeds.

4.4. Value of sharing goals

We use our method to evaluate the value of satellites sharing their goals (i.e., sharing their orbits, and any associated changes or maneuvers) with each other. In our experiments, a randomly chosen satellite changes its goal midway through the episode. The new goal is set to a uniformly-random location that is within a distance ρ_{\max} of the original goal, as illustrated in Figure 5.




Figure 5: For one of the three agents, the goal is randomly resets midway through the episode to be within a distance ρ_{\max} from p_i^{goal} .




Figure 6: Percentage improvement achieved through goal-sharing for (1) the success rates, S (in blue; left-axis), and (2) the fraction of episode (or time) taken on average by agents to reach their goal, T (in red; right-axis). Moving averages over 0.2 km increases in ρ_{\max} are shown.

We progressively increase ρ_{\max} in steps of 20 m from zero to 1 km (which considering the 2-km \times 2-km environment size is approximately equivalent to resetting the goal randomly somewhere within the environment). We generate 100 instances for each value of ρ_{\max} , with random initializations of satellite locations and their goals. For each such instance, we use InforMARL trained on the space environment to evaluate the cases when the satellites share their goals with each other, and when they do not. We consider two performance measures from Section 4.1, the success rate (S) and the average time to reach goal (T). We evaluate the performance improvement from goal-sharing as follows: For the success rates, we calculate $\frac{S_{\text{goal_sharing}} - S_{\text{no_goal_sharing}}}{S_{\text{no_goal_sharing}}}$ as a percentage. The improvement in the average time to reach the goal is calculated similarly, accounting for the fact that an improvement corresponds to a decrease in T .

Figure 6 demonstrates the performance improvement (relative to the performance without goal sharing) that is achieved through goal sharing, as the maximum goal reset distance increases. We note that positive values in Figure 6 indicate that the success rates *increase* with goal-sharing and the times taken by agents to reach their goals *decrease*, illustrating the benefits of goal-sharing for all values of ρ_{\max} considered. The median number of collisions was found to be the same irrespective of goal-sharing, so we do not discuss this performance metric further.

For small values of ρ_{\max} (e.g., $\rho_{\max} < 0.3$ km in Figure 6), the approximate location of the new goal can be inferred from the current goals/orbits, so we see relatively modest performance improvement in success rates from goal-sharing. As ρ_{\max} increased further, we see significant improvements (reaching a more than 100% increase) in the success rates from goal-sharing. Interestingly, we see that for higher values of ρ_{\max} (e.g., $\rho_{\max} > 0.65$ km in Figure 6), the benefits of goal-sharing begin to decrease. We hypothesize that this behavior is related to relationship between the aleatory uncertainty of satellite trajectories and collision risk (Balch et al., 2019). Let us consider the example described by Balch et al. (2019): Suppose two satellites are known with certainty to be on a collision course (i.e., there is no epistemic uncertainty). If we knew that one of the operators could apply an arbitrarily large maneuver in a random, unknown direction, then the probability of a collision post-maneuver decreases as the magnitude of the possible maneuver increases. Analogously in our situation, this would mean that for large enough values of ρ_{\max} , the success rates start increasing even without goal-sharing; consequently, the percentage improvement in these rates obtained from goal sharing starts decreasing as seen in Figure 6.

In summary, we find that goal-sharing generally improves success rates and decreases times to goal, demonstrating the potential benefits of such information-sharing among satellite operators.

5. Conclusions and Future Work

We demonstrated that transfer learning from a ground-based environment to a space-based one can improve both sample complexity and performance, despite the significant differences in the underlying dynamics that govern the agents in the two environments. We also found that InforMARL, our GNN-based approach, is scalable in a space-based environment, satisfying a critical need for space traffic management as the skies become more dense with satellites and debris. Our initial investigations showed that goal sharing among satellite operators can improve the safety and efficiency of space traffic operations. Future work will include developing a more realistic space traffic simulation environment, accounting for communication delays and losses, adding mechanisms to provide safety guarantees, and designing incentive mechanisms for information-sharing among satellite operators.

Acknowledgments

The authors would like to thank the MIT SuperCloud (Reuther et al., 2018) and the Lincoln Laboratory Supercomputing Center for providing high performance computing resources that have contributed to the research results reported within this paper. We also thank Jasmine Aloor and Victor Qin for their help running some of the simulations. The NASA University Leadership Initiative (grant #80NSSC20M0163) provided funds to assist the authors with their research, but this article solely reflects the opinions and conclusions of its authors and not any NASA entity. This research was sponsored in part by the United States AFRL and the United States Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notion herein. Sydney Dolan was supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1650114.

References

- Johannes Ackermann, Volker Gabler, Takayuki Osa, and Masashi Sugiyama. Reducing overestimation bias in multi-agent domains using double centralized critics. *CoRR*, abs/1910.01465, 2019. URL <http://arxiv.org/abs/1910.01465>.
- Akshat Agarwal, Sumit Kumar, and Katia P. Sycara. Learning transferable cooperative behavior in multi-agent teams. *CoRR*, abs/1906.01202, 2019. URL <http://arxiv.org/abs/1906.01202>.
- Michael Scott Balch, Ryan Martin, and Scott Ferson. Satellite conjunction analysis and the false confidence theorem. *Proceedings of the Royal Society A*, 475(2227), July 2019. URL <http://doi.org/10.1098/rspa.2018.0565>.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Christopher Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique Pondé de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. Dota 2 with large scale deep reinforcement learning. *CoRR*, abs/1912.06680, 2019. URL <http://arxiv.org/abs/1912.06680>.
- Claudio Bombardelli. Analytical formulation of impulsive collision avoidance dynamics. *Celestial Mechanics and Dynamical Astronomy*, 118(2):99–114, 2013. doi: 10.1007/s10569-013-9526-3.
- Claudio Bombardelli and Javier Hernando-Ayuso. Optimal impulsive collision avoidance in low earth orbit. *Journal of Guidance, Control, and Dynamics*, 38(2):217–225, 2015. doi: 10.2514/1.g000742.
- Lin Cheng, Zhenbo Wang, and Fanghua Jiang. Real-time control for fuel-optimal moon landing based on an interactive deep reinforcement learning algorithm. *Astrodynamic*s, 3(4):375–386, 2019. doi: 10.1007/s42064-018-0052-2.

- Richard Cheng, Mohammad Javad Khojasteh, Aaron D. Ames, and Joel W. Burdick. Safe multi-agent interaction through robust control barrier functions with learned uncertainties. *CoRR*, abs/2004.05273, 2020. URL <https://arxiv.org/abs/2004.05273>.
- Andrea De Vittori, Maria Francesca Palermo, Pierluigi Di Lizia, and Roberto Armellin. Low-thrust collision avoidance maneuver optimization. *Journal of Guidance, Control, and Dynamics*, 45(10): 1815–1829, 2022. doi: 10.2514/1.g006630.
- Jeff Foust. <https://spacenews.com/spacex-and-oneweb-spar-over-satellite-close-approach/>. *SpaceNews*, Apr 2021. URL <https://spacenews.com/spacex-and-oneweb-spar-over-satellite-close-approach/>.
- Niko A. Grupen, Bart Selman, and Daniel D. Lee. Fairness for cooperative multi-agent learning with equivariant policies, 2021.
- Ryan Hiles, Amos Alexander, Nicholas Herzer, Diana McKissock, Aaron Mitchell, and Jonathan Sapinosa. Report on 2020 mega-constellation deployments and impacts to space domain awareness. AMOS, 2021. URL <https://amostech.com/TechnicalPapers/2021/SSA-SDA/Hiles.pdf>.
- Heather Hinkel, Scott P. Cryan, and Chris D’Souza. *Rendezvous and Docking Strategy for Crewed Segment of the Asteroid Redirect Mission*. doi: 10.2514/6.2014-1796. URL <https://arc.aiaa.org/doi/abs/10.2514/6.2014-1796>.
- Kerianne Hobbs, Alexander R. Collins, and Eric Feron. Towards a taxonomy for automatic and autonomous cooperative spacecraft maneuvering in a space traffic management framework. *ASCEND 2020*, 2020. doi: 10.2514/6.2020-4240.
- Kirk Hovell and Steve Ulrich. On deep reinforcement learning for spacecraft guidance. *AIAA Scitech 2020 Forum*, 2020. doi: 10.2514/6.2020-1600.
- Dario Izzo, Marcus Märtens, and Binfeng Pan. A survey on artificial intelligence trends in spacecraft guidance dynamics and control. *Astrodynamics*, 3(4):287–299, 2019. doi: 10.1007/s42064-018-0053-6.
- Arbaaz Khan, Ekaterina Tolstaya, Alejandro Ribeiro, and Vijay Kumar. Graph policy gradients for large scale robot control, 2019.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *CoRR*, abs/1706.02275, 2017a. URL <http://arxiv.org/abs/1706.02275>.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *Neural Information Processing Systems (NIPS)*, 2017b. URL <https://github.com/openai/multiagent-particle-envs>.
- Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2):1–141, 2017. doi: 10.1561/22000000060. URL <https://doi.org/10.1561%2F22000000060>.

- Siddharth Nayak, Kenneth Choi, Wenqi Ding, Sydney Dolan, Karthik Gopalakrishnan, and Hamsa Balakrishnan. Scalable multi-agent reinforcement learning through intelligent information aggregation. In *International Conference on Machine Learning*, 2023.
- John O’Callaghan. Hope for solving space debris. *Aerospace America*, 2022.
- V.G. Rao and D.S. Bernstein. Naive control of the double integrator. *IEEE Control Systems Magazine*, 21(5):86–97, 2001. doi: 10.1109/37.954521.
- Tabish Rashid, Mikayel Samvelyan, Christian Schröder de Witt, Gregory Farquhar, Jakob N. Foerster, and Shimon Whiteson. QMIX: monotonic value function factorisation for deep multi-agent reinforcement learning. *CoRR*, abs/1803.11485, 2018. URL <http://arxiv.org/abs/1803.11485>.
- James Rendleman and Sarah Mountin. Responsible ssa cooperation to mitigate on-orbit space debris risks. In *International Conference on Recent Advances in Space Technologies*, page 851–856, 2015.
- Albert Reuther, Jeremy Kepner, Chansup Byun, Siddharth Samsi, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Anna Klein, Lauren Milechin, Julia Mullen, Andrew Prout, Antonio Rosa, Charles Yee, and Peter Michaleas. Interactive supercomputing on 40,000 cores for machine learning and data analysis. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2018.
- Jennifer A Roberts and Peter C. E. Roberts. The development of high fidelity linearised j2 models for satellite formation flying control. In *14th AAS/AIAA Space Flight Mechanics Meeting*, Feb 2004.
- G Salemme, R Armellin, and P Di Lizia. *AIAA SciTech*. Jan 2020. URL <https://arc.aiaa.org/doi/pdf/10.2514/6.2020-0231>.
- Peng Mun Siew, Daniel Jang, Thomas G. Roberts, and Richard Linares. Space-based sensor tasking using deep reinforcement learning. *Journal of the Astronautical Sciences*, 2022. doi: <https://doi.org/10.1007/s40295-022-00354-8>.
- Space Debris Office. *ESA’S Annual Space Environment Report*. ESA, 2022. URL https://www.sdo.esoc.esa.int/environment_report/Space_Environment_Report_latest.pdf.
- Christopher Sullivan and Natasha Bosanac. Using multi-objective deep reinforcement learning to uncover a pareto front in multi-body trajectory design. In *AAS/AIAA Astrodynamics Specialist Conference*, 2020.
- Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinícius Flores Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z. Leibo, Karl Tuyls, and Thore Graepel. Value-decomposition networks for cooperative multi-agent learning. *CoRR*, abs/1706.05296, 2017. URL <http://arxiv.org/abs/1706.05296>.
- David A. Vallado and Wayne D. McClain. *Fundamentals of astrodynamics and applications*. Microcosm Press, 2007.

Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, L. Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander Sasha Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom Le Paine, Caglar Gulcehre, Ziyun Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy P. Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, pages 1–5, 2019.

Ulrich Walter. Orbit perturbations. *Astronautics*, page 555–660, 2018. doi: 10.1007/978-3-319-74373-8_12.

Charity Weeden, Jonathan Goff, Jessica Noble, Jeremy Schiel, Dana Turse, Colin Doughan, John Carrico, and Rob Patterson. Comments of global newspace operators, 2019. URL <https://www.fcc.gov/ecfs/search/search-filings/filing/1040578949828>.

Tianpei Yang, Weixun Wang, Hongyao Tang, Jianye Hao, Zhaopeng Meng, Hangyu Mao, Dong Li, Wulong Liu, Chengwei Zhang, Yujing Hu, Yingfeng Chen, and Changjie Fan. An efficient transfer learning framework for multiagent reinforcement learning, 2020. URL <https://arxiv.org/abs/2002.08030>.

Chao Yu, Akash Velu, Eugene Vinyals, Jiayuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of PPO in cooperative multi-agent games. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. URL <https://openreview.net/forum?id=YVXaxB6L2P1>.