

# **Multi-Robot Formation Control Using Reinforcement Learning**

by

Abhay Rawat, Kamalakar Karlapalem

Report No: IIIT/TR/2020/-1



Centre for Data Engineering  
International Institute of Information Technology  
Hyderabad - 500 032, INDIA  
January 2020

# Multi-Robot Formation Control Using Reinforcement Learning

Abhay Rawat and Kamalakar Karlapalem

**Abstract**—In this paper, we present a machine learning approach to move a group of robots in a formation. We model the problem as a multi-agent reinforcement learning problem. Our aim is to design a control policy for maintaining a desired formation among a number of agents (robots) while moving towards a desired goal. This is achieved by training our agents to track two agents of the group and maintain the formation with respect to those agents. We consider all agents to be homogeneous and model them as unicycle [1]. In contrast to the leader-follower approach, where each agent has an independent goal, our approach aims to train the agents to be cooperative and work towards the common goal. Our motivation to use this method is to make a fully decentralized multi-agent formation system and scalable for a number of agents.

## I. INTRODUCTION


Multi-robot system has recently become a popular problem in robotics. Some of the work done on multi-robot and multi agent systems are area exploration [2] [3], payload transportation [4] [5], etc. One of the classic examples of multi-robot coordination is formation control of agents, i.e. a group of agents moving while maintaining a particular shape. Various control algorithms have been implemented to move the agents in formation which are typically based on leader-follower control architecture [6], virtual structure [7] control architecture and behaviour modeling based control [8].

Among the above mentioned architectures, leader-follower based control architecture is the most popular due to its easy implementation and robustness. In leader-follower based approach, a leader is assigned to the formation that moves towards the desired goal using desired motion planning and tracking algorithms. All the other agents maintain a specific distance and angle with respect to the leader robot to maintain the formation.

In our work, we explore the formation control problem through a coordination based approach using reinforcement learning. *Instead of following a designated leader, the agents learn to coordinate with each other to achieve the common goal (go to goal while maintaining a given formation).*

We also propose to make this method scalable with respect to the number of agents in formation and also the shape of the formation. The policies of the agents are trained to track any two (or any predefined number) of their neighbours. Once the policies are trained, this model can be used for different number of agents and formation shapes.

In [9] [10], the authors use reinforcement learning for formation control in conjunction with leader-follower approach. The agent (robot) would learn how to maintain a specified distance from a designated leader. In [11], the authors use reinforcement based control alongside a behaviour based formation controller. The actions of the agent would be



**Fig. 1:** A multi exposure image of three agents(denoted as triangles) moving in an equilateral triangle formation towards the goal

the average of the control signals of both the controllers. In both of the above mentioned cases, the agents work towards individual goals to complete the task, whereas in our work we aim to develop coordination between agents to achieve that common goal. Unlike [12] and [13], we do not consider communication amongst the agents and investigate the effectiveness of our method without it.

## II. PROBLEM FORMULATION

We formulate the task of distributed formation control as a discounted stochastic game  $G$ , described by the tuple  $G = \langle S, U, A, P, R, Z, O, \gamma \rangle$ . The true state of the environment is denoted by  $s \in S$ , where  $S$  denotes the state space of the environment. At each time-step every agent  $a \in A$  takes an action  $u_a \in U$  resulting in the transition of the environment to some other state  $s'$ . The transition is stochastic, and the probability of the next state  $s'$  conditioned on the actions  $u$  and the state  $s$  is given by  $P(s'|s, \bar{u}) : S \times U \times S \rightarrow [0, 1]$ , where  $\bar{u} \in U$  is the joint action of all the agents and  $U = U^{|A|}$  is the joint action space. Due to the cooperative nature of the task all the agents share a common reward function  $R(s, \bar{u}) : S \times U \rightarrow \mathbb{R}$  and  $\gamma \in [0, 1]$  is the discount factor. At any given time-step, an agent  $a$  observes the partial state of the environment  $o \in O$  according to the observation function  $Z(s, a) : S \times A \rightarrow O$ .

Each agent interacts with the environment according to its own policy. Each agent has its history of interactions with the environment which it uses to train its policy. Following the work of the authors from [14] [15], we use the paradigm of centralized training with decentralized execution. We use

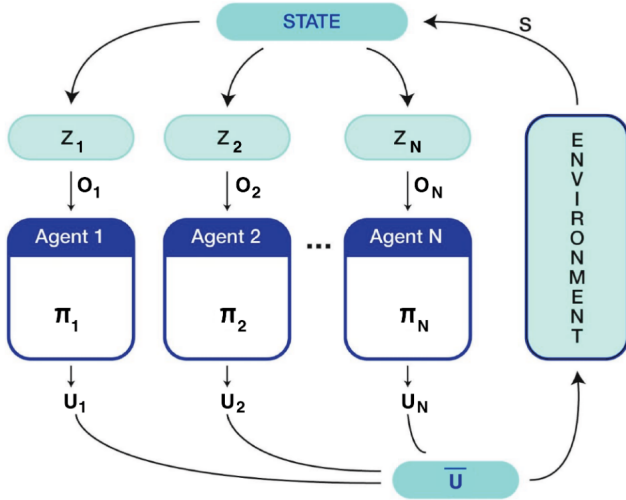
the actor-critic architecture with a central critic that aims to approximate the state-action value function. The central critic would represent a state action value function which maps the true state of the environment and the joint actions, to the expected sum of rewards. We use a deep neural network  $Q$  parameterized by  $\theta_q$  to approximate this function. In particular, the critic function aims to minimize the following loss:

$$J_Q(s, \bar{u}) = \|\hat{y} - Q(s, \bar{u}|\theta_q)\|_2 \quad (1)$$

where  $\hat{y} = R(s, \bar{u}) + \gamma \max_{\bar{u}'} \bar{Q}(s', \bar{u}')$ ,  $\bar{u}'$  is the next joint action taken by the agents according to their current policy, and  $\bar{Q}$  is the target network (parameterized by  $\bar{\theta}_q$ ) as used by [16]. The parameters of the target critic network are updated by taking a weighted average of the parameters of both the critic networks (polyak averaging [17]). The weights of the target critic network are updated as

$$\bar{\theta}_q = \tau \theta_q + (1 - \tau) \bar{\theta}_q \quad (2)$$

where  $\tau \ll 1$ . This update (soft update) has been found to stabilize the learning process [18] by reducing the variance in the target  $Q$  values used to estimate the loss in Eq. 1.



**Fig. 2:** Interaction of agents with the environment. The state of the environment is observed through the observation function  $Z_i$  of the agent  $i$ . The agent acts according to its policy  $\pi_i$ . The actions of all the agents  $\bar{u}$ , transitions the environment to some new state.

We use policy gradient method [19] [18] to train the actor network  $\pi$ , parameterized by  $\theta_\pi$ , by adjusting its parameters ( $\theta_\pi$ ) to maximize the expected reward  $J_\pi$ , where

$$J_\pi(\cdot, \theta_\pi) = \mathbb{E}_{s \sim p^\pi, a \sim \pi} [R(s, \bar{u})]$$

Every agent's policy is trained by updating its parameters in the direction of  $\nabla_{\theta_i} J_\pi$  (gradient of the loss  $J_\pi$  with respect to the parameters of  $\pi_i$ ). Policy gradient methods are known to result in high variance estimates of the gradient. In multi-agent setting this is quite an issue as the agents gradients are also dependent on the action of other agents [15]. To mitigate this problem to some extent we stick to the deterministic flavour of the algorithm proposed in [18].

For distributed formation control, we leverage the notion of minimal rigidity and persistent graphs [20] to define our system. We denote each agent as a vertex  $v_i \in V$  in a graph  $\mathbb{G} = (V, E)$ . There is a directed edge  $e_{i,j} = (v_i, v_j) \in E$  between two agents  $v_i, v_j$  if and only if agent  $v_i$  observes agent  $v_j$ . Following [20], we use the fact that if each of the agent in the formation observes at least two other agents then we have a minimally rigid graph. We can use these trained agents for any shape of formation by providing them the inter-agent distances. In addition to observing other agents, all the agents continuously track the goal position. The agents are rewarded if they are able to maintain the specific distance with the agents that they can observe. A reward of greater magnitude is given if the agents are able to reach the goal while maintaining said formation. To define the reward function, we mathematically introduce the following conditions for judging actions at the given states.

[C1] **Formation condition:** If all the agents maintain their specified edge lengths within an error of  $\epsilon_{form}$

$$|d_{i,j}(s) - \bar{d}_{i,j}| \leq \epsilon_{form} \quad \forall i, j \in E \quad (3)$$

[C2] **Collision condition:** This condition is met if the distance between any two agents is less than some threshold (i.e., collision).  $\epsilon_{col}$ .

$$|d_{i,j}(s)| \leq \epsilon_{coll} \quad \forall i, j \in E \quad (4)$$

[C3] **Success condition:** If formation condition is satisfied and

$$|c_V - g| \leq \epsilon_{goal} \quad (5)$$

where  $d_{i,j}(s)$  represents the distance between agents  $i$  and  $j$  at state  $s$ ,  $\bar{d}_{i,j}$  denotes the desired distance between agents  $i$  and  $j$ ,  $c_V$  denotes the centroid of the formation made by the agents in  $V$ ,  $g$  is the desired goal point of the formation,  $\epsilon_{form}$  and  $\epsilon_{goal}$  are the tolerances in error. The reward function is defined as:

$$R(s, a, s') = \begin{cases} r_{edge}, & \text{if only [C1] is true} \\ r_{collision}, & \text{if [C2] is true} \\ r_{goal}, & \text{if [C3] is true} \\ r_{penalty}, & \text{otherwise} \end{cases} \quad (6)$$

To encourage the agents to complete the task as soon as possible, we give a small negative penalty  $r_{penalty}$ . We also ensure that  $r_{edge} \ll r_{goal}$  so that the agents do not get stuck in a local minima (where they remain in formation while totally ignoring the high level goal). The values for these parameters that were used in our experiments are described in Table I.

### III. METHODOLOGY

The control inputs of each agent is its linear and angular velocity ( $v$  and  $\omega$ ). Each agent  $i$  observes up to two other agents along with the target location. The observations are made in the local reference frames of the observing agents. Each agent has its own independent policy  $\pi$  which maps the observation space to the action space. We use a neural

**TABLE I:** List of Parameters.

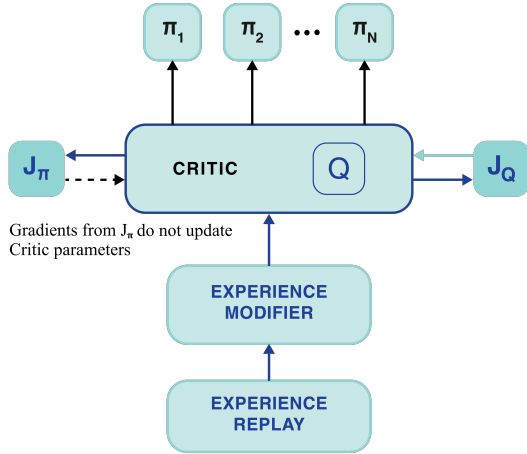
Parameter Name	Value
H	15
$\epsilon_{form}$	0.10 m
$\epsilon_{goal}$	0.15 m
$\epsilon_{coll}$	0.20 m
$r_{edge}$	0.1
$r_{collision}$	-100
$r_{goal}$	50
$r_{penalty}$	-0.5
$\tau$	5e-3

network to define this function with parameters  $\theta_\pi$ . To mitigate the problem of partial observability we provide the agents a sequence of observations of length  $H$ . This helps the agents to learn the behaviour of other agents. In particular, the actor function (policy) can be described as:

$$\pi(o_{t-H-1}, \dots, o_t) : O^H \rightarrow U$$

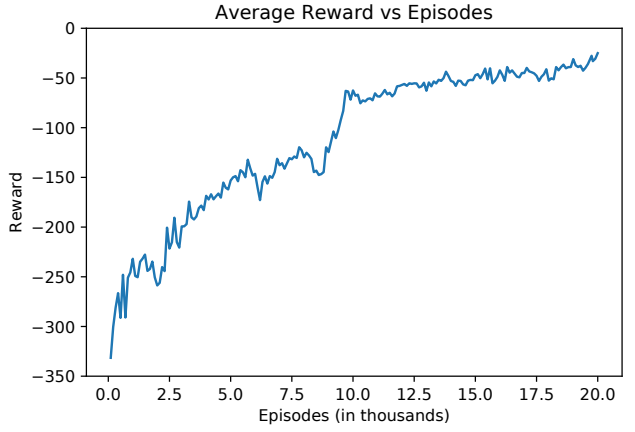
We define a central critic [14] that maps the true state of the system and the joint actions taken by the agents at that state to the expected future return (state-action value). The true state of the system is defined as the pose information of all the agents with respect to the centroid of the current formation i.e.  $[x_i^c, y_i^c, \theta_i^c] \forall i \in V$  augmented with the position vector of the goal state with respect to the centroid. The critic function is described as:

$$Q(s_{t-H-1}, \dots, s_t, \bar{u}_{t-H-1}, \dots, \bar{u}_t) : S^H \times U^H \rightarrow U$$



**Fig. 3:** The training involves sampling a batch of experiences, then modifying some of the experiences for positive reinforcement.

We build a centralized experience replay containing tuples of the past experiences  $\langle s, r, s', u_1, \dots, u_{|A|}, o_1, \dots, o_{|A|}, o'_1, \dots, o'_{|A|} \rangle$ , where  $s'$  and  $o'$  are the next state of the environment and the observation made by the agent  $i$  for the corresponding state  $s'$  respectively. At each training iteration we randomly sample a sequence of  $H$  concurrent experiences from this memory. We use a similar approach as proposed by the authors of [21], to sample concurrent experiences.



**Fig. 4:** Average reward per episode vs. number of training episodes

To speed up learning, we modify the batch of experiences using a similar method as in [22]. This is shown as experience modifier in Fig. 3. The authors in [22] change the goal state to the actual state achieved by the agent. This leads in more positive reinforcement and thus speeds up the learning process considerably. In this case, the goal is observed with respect to different observation frames and the input to the networks is a sequence containing those goals at different time-steps. In order to make the changes consistent throughout the sequence, we keep track of the transformations from previous pose to the current pose, for every agent. For a sequence of  $H$  observations  $o_1, \dots, o_H$ , let the goal achieved be  $g_t$ , and the transformation matrix between the consecutive poses of an agent be  $K_t^{t+1}$ , then:

$$g_t = K_t^{t+1} \times g_{t+1} \quad t \in 1, \dots, H-1 \quad (7)$$

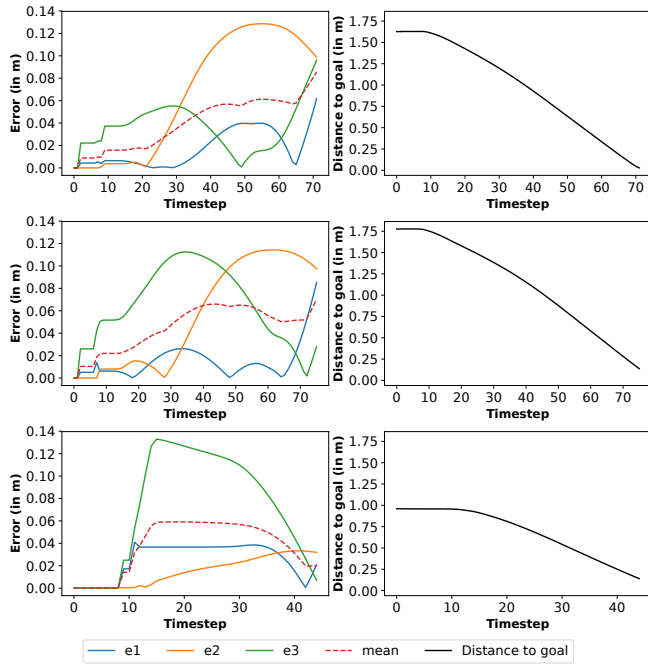
where  $g_H$  is the centroid of the formation at time-step  $H$ .

#### IV. RESULTS

The simulations are performed using three agents. The agents are given a desired formation to maintain and a desired goal. During execution, only the actor network is used, which can be easily executed on low cost mobile computers such as Raspberry Pi.

Fig. 4 shows the learning curve for training. The flow of the training is described in Fig. 3. We trained the agents for 20000 episodes, where each episode had a maximum of 120 time-steps. An episode could terminate early if there is any collision amongst the agents, or task was completed within 120 time-steps. The environment is a custom environment made with the help of OpenAI gym [23]. The arena is a square of side length 10 m. During training, the goal state, the centroid of the formation and the formation to be maintained were generated at random. To facilitate the training process, The initial states of individual agents were computed from the centroid and the target formation to make sure that the agents were in formation in the initial state.

Fig. 5 shows the error in each edge of the formation alongside the distance of the centroid from the goal with respect to time for the corresponding trial.



**Fig. 5:** The plots on the left show the error in each edge of the formation. The plots on the right show the distance between the centroid of the formation and the goal. The edges between agents are denoted by  $e1$ ,  $e2$  and  $e3$ . The dotted red line denotes the mean error of all the edges.

## V. CONCLUSION

In contrast to existing formation control algorithms, we aim to build a scalable and modular control mechanism for formations through multi-agent reinforcement learning.

Multi-agent systems suffer from non-stationarity of the environment with respect to the agents. This can cause fluctuations in the learning process as well as during execution. In a task like formation control, this can result in cascading effect and the agents might not be able to recover from a poor choice of action. This effect can be seen in the first two rows of Fig. 5, where the error in the edges fluctuates throughout the episode. In future work we aim to improve the accuracy of the formation and to address the problem of non-stationarity which hinders the performance when more agents are added. To mitigate this issue we can limit the change in policy at every iteration for each agent. Such method has already been used by [24] [25] to prevent the policy to change drastically. By doing this we hope that the agents would be able to learn about the policies of other agents and take actions accordingly.

## REFERENCES

- [1] R. C. Arkin, R. C. Arkin, *et al.*, *Behavior-based robotics*. MIT press, 1998.
- [2] A. Baranzadeh and A. V. Savkin, "A distributed control algorithm for area search by a multi-robot team," *Robotica*, vol. 35, no. 6, pp. 1452–1472, 2017.
- [3] S. Sharma, A. Shukla, and R. Tiwari, "Multi robot area exploration using nature inspired algorithm," *Biologically Inspired Cognitive Architectures*, vol. 18, pp. 80–94, 2016.

- [4] P. Verma, R. Tallamraju, A. Rawat, S. Chand, and K. Karlapalem, "Loosely coupled payload transport system with robot replacement," *arXiv preprint arXiv:1904.03049*, 2019.
- [5] Y. Sirineni, P. Verma, and K. Karlapalem, "Traffic management strategies for multi-robotic rigid payload transport systems," *arXiv preprint arXiv:1906.11452*, 2019.
- [6] Z. Peng, G. Wen, A. Rahmani, and Y. Yu, "Leader-follower formation control of nonholonomic mobile robots based on a bioinspired neurodynamic based approach," *Robotics and autonomous systems*, vol. 61, no. 9, pp. 988–996, 2013.
- [7] C. B. Low and Q. San Ng, "A flexible virtual structure formation keeping control for fixed-wing uavs," in *2011 9th IEEE international conference on control and automation (ICCA)*, pp. 621–626, IEEE, 2011.
- [8] T. Balch and R. C. Arkin, "Behavior-based formation control for multirobot teams," *IEEE transactions on robotics and automation*, vol. 14, no. 6, pp. 926–939, 1998.
- [9] M. Knopp, C. Aykin, J. Feldmaier, and H. Shen, "Formation control using q ( $\lambda$ ) reinforcement learning," in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 1043–1048, IEEE, 2017.
- [10] G. Zuo, J. Han, and G. Han, "Multi-robot formation control using reinforcement learning method," in *International Conference in Swarm Intelligence*, pp. 667–674, Springer, 2010.
- [11] R. J. Johns, "Intelligent formation control using deep reinforcement learning," 2018.
- [12] M. Turpin, N. Michael, and V. Kumar, "Decentralized formation control with variable shapes for aerial robots," in *2012 IEEE international conference on robotics and automation*, pp. 23–30, IEEE, 2012.
- [13] T.-H. Cheng, Z. Kan, J. A. Rosenfeld, and W. E. Dixon, "Decentralized formation control with connectivity maintenance and collision avoidance under limited and intermittent sensing," in *2014 American nuclear conference*, pp. 3201–3206, IEEE, 2014.
- [14] J. N. Foerster, G. Farquhar, T. Afouras, N. Nardelli, and S. Whiteson, "Counterfactual multi-agent policy gradients," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [15] R. Lowe, Y. Wu, A. Tamar, J. Harb, O. P. Abbeel, and I. Mor-datch, "Multi-agent actor-critic for mixed cooperative-competitive environments," in *Advances in Neural Information Processing Systems*, pp. 6379–6390, 2017.
- [16] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double q-learning," in *Thirtieth AAAI conference on artificial intelligence*, 2016.
- [17] B. T. Polyak and A. B. Juditsky, "Acceleration of stochastic approximation by averaging," *SIAM Journal on Control and Optimization*, vol. 30, no. 4, pp. 838–855, 1992.
- [18] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [19] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Ried-miller, "Deterministic policy gradient algorithms," 2014.
- [20] J. M. Hendrickx, B. D. Anderson, and V. D. Blondel, "Rigidity and persistence of directed graphs," in *Proceedings of the 44th IEEE Conference on Decision and Control*, pp. 2176–2181, IEEE, 2005.
- [21] S. Omidshafiei, J. Pazis, C. Amato, J. P. How, and J. Vian, "Deep decentralized multi-task multi-agent reinforcement learning under partial observability," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2681–2690, JMLR. org, 2017.
- [22] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welin-der, B. McGrew, J. Tobin, O. P. Abbeel, and W. Zaremba, "Hindsight experience replay," in *Advances in Neural Information Processing Systems*, pp. 5048–5058, 2017.
- [23] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [24] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [25] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, "Trust region policy optimization," *CoRR*, vol. abs/1502.05477, 2015.