

# Learning Neural Traffic Rules

Xuan Zhang, Xifeng Gao, Kui Wu, Zherong Pan<sup>†</sup>

**Abstract**—Extensive research has been devoted to the field of multi-agent navigation. Recently, there has been remarkable progress attributed to the emergence of learning-based techniques with substantially elevated intelligence and realism. Nonetheless, prevailing learned models face limitations in terms of scalability and effectiveness, primarily due to their agent-centric nature, i.e., the learned neural policy is individually deployed on each agent. Inspired by the efficiency observed in real-world traffic networks, we present an environment-centric navigation policy. Our method learns a set of traffic rules to coordinate a vast group of unintelligent agents that possess only basic collision-avoidance capabilities. Our method segments the environment into distinct blocks and parameterizes the traffic rule using a Graph Recurrent Neural Network (GRNN) over the block network. Each GRNN node is trained to modulate the velocities of agents as they traverse through. Using either Imitation Learning (IL) or Reinforcement Learning (RL) schemes, we demonstrate the efficacy of our neural traffic rules in resolving agent congestion, closely resembling real-world traffic regulations. Our method handles up to 240 agents at real-time and generalizes across diverse agent and environment configurations.

**Index Terms**—Traffic Rule, Multi-Agent Navigation, Neural Navigation Policy

## I. INTRODUCTION

Multi-agent navigation forms the cornerstone of numerous pivotal robotic applications, including domains such as automated warehousing [1], autonomous driving [2], and realization of smart cities [3]. As a consequence, the refinement of navigation algorithms has garnered substantial research focus throughout the past decades. An ideal navigation algorithm should satisfy three desiderata. Scalability: Capable of controlling an arbitrarily large crowd of agents; Generality: Capable of handling arbitrary environment and agent configurations; Efficacy: A mild growth of computational and deployment cost with crowd size. Unfortunately, to this day, an ideal navigation algorithm that satisfies all three attributes continues to be elusive.

The key to a successful navigation algorithm is the strategy to mitigate agent congestion. Early research endeavors tackled congestion through either back-tracking search [4, 5, 6] or localized collision-avoidance [7, 8]. The search-based algorithms exhibit notable generality; some even provide assurances of completeness or optimality. However, they have limitations with scalability and efficacy. Conversely, the local navigation policies exhibit ideal scalability, yet their generality is curtailed due to their myopic nature, oftentimes leading to agents stuck in congested configurations. Recent strides in this domain have been achieved through learning-based navigation policies [9, 10, 11]. These innovative algorithms parameterize navigation policies as deep networks, subsequently fine-tuned through either IL or RL to address complex, long-term decision-making challenges. While learning-based policies empirically

surpass myopic local navigation methods, they necessitate that each agent possesses the computational resources to execute intricate neural network inferences. This assumption results in a substantial deployment cost. In reality, large-scale robot swarm systems [12, 13] are often equipped with limited computational capabilities that cannot feasibly accommodate the demands of neural network inferences.

Examining prevailing learning-based policies [9, 14, 10], we find a vast majority of them being agent-centric, i.e., their deep neural policies are independently deployed on each individual agent, which inevitably requires non-trivial computation and communication capabilities. In contrast, real-world traffic networks exhibit an environment-centric approach rather than an agent-centric one. At the heart of these networks resides a collection of human-designed traffic rules driven solely by the specific characteristics of the environment, such as highways, crossroads, and T-junctions. The efficacy of these rules is evidenced in their ability to manage substantial crowds of pedestrians and vehicles, effortlessly mitigating congestion by adhering to predefined traffic regulations. In essence, this approach sidesteps the need for intricate agent-specific decision-making processes.

**Main Results:** Inspired by the elegance of real-world traffic networks, we introduce a learnable environment-centric navigation policy, assuming that our agents are endowed solely with rudimentary computational capacities, confining their interactions to local navigations. Our neural policy encapsulates a set of traffic rules based on the environments, which are followed by our agents. Specifically, our method decomposes the environment into discrete blocks, and we represent our neural policy as a GRNN built on the block network. Each GRNN node is trained to manipulate the velocities of agents traversing the corresponding block, facilitating the rule-following behavior. Our GRNN is trained to mitigate agent congestion via two methodologies: IL, guided by a groundtruth traffic rule, and RL, guided by a congestion-resolving reward signal. After the training phase, our neural policy can be deployed on unseen environment configurations to coordinate up to 240 intelligent agents at real-time in a simulated environment. Our contributions are summarized below:

- A decentralized agent navigation paradigm utilizing learned environment-encoded traffic rules.
- Environment-centric navigation policies parameterization using GRNN.
- Reward design and training algorithms for environment-centric policies in both IL and RL setting.

## II. RELATED WORK

We review related works on multi-agent navigation, learning-based navigation methods, and traffic-rule modeling.

**Multi-Agent Navigation:** Over the course of time, researchers have branched into two avenues, yielding centralized

<sup>†</sup> indicates corresponding author. Authors are with Lightspeed Studios, Tencent America. {xuanzzhang,xifengao,kuiwu,zrpan}@global.tencent.com.

and decentralized navigation algorithms. Centralized algorithms operate under the premise that a central computational node coordinates the movements of all agents. This node employs various search-based algorithms to avert congestion over a long horizon. One noteworthy example of this paradigm is the conflict-based search algorithm [6]. It is a variant of the branch-and-bound approach, designed to resolve motion conflicts among agents iteratively converging to the optimal navigation plan. While such methods theoretically accommodate any environment or agent configuration, their practical computational demands surge with the number of agents, as analyzed in [5]. Consequently, real-world centralized algorithms either compromise optimality for efficiency [15] or lean on additional assumptions [16]. Despite their comprehensive applicability, centralized algorithms exert a substantial communication cost, necessitating synchronization among all agents in each step.

Considering their limitations, research endeavors have veered towards decentralized algorithms. In this context, each agent autonomously plans and executes its motion asynchronously. Prominent decentralized techniques [7, 8] only mandate agents to engage in rudimentary collision avoidance with local neighbors, devoid of extensive long-horizon decision-making. Astonishingly, these straightforward algorithms capably resolve congestion for sizable agent swarms in environments replete with expansive open spaces. Nonetheless, in intricate settings with narrow passages, agents can become ensnared. Such deadlock configurations can be alleviated by introducing heuristic congestion-resolving behaviors [17], such as grouping [18], yielding [19], and local coordination [20]. Regrettably, while these heuristic approaches ameliorate specific types of deadlock scenarios, none possess the potential to address all scenarios comprehensively. In contrast, our proposed method presents a cohesive strategy for tackling congestion issues while preserving the efficiency and scalability inherent in decentralized algorithms. Our neural policy embeds environment-centric traffic regulations and operates under the assumption of minimal computational capabilities on the agents' part.

**Learning-based Navigation Methods:** The incorporation of data-driven techniques has proven highly effective in enhancing the performance of navigation algorithms. In the realm of centralized algorithms, Huang *et al.* [21] introduced a method to learn a conflict selection policy, thereby expediting the conflict-based search process. Another significant advancement comes from the work of Han and Yu [22], who harnessed a database of small-scale navigation plans to accelerate the search for larger-scale counterparts. Most other data-driven approaches are designed to work with decentralized algorithms. The challenge of addressing deadlock concerns in local navigation [7, 8] prompted Fan *et al.* [9] to propose a neural navigation policy, trained through RL, for mitigating long horizon congestion. However, their neural policy's reliance solely on local environmental information precludes it from surpassing the performance of conventional local navigation algorithms.

Enhancements to neural policies have subsequently emerged, including the use of global environment maps [23] and the facilitation of inter-agent communication [14, 24].

Notably, these approaches, driven by an agent-centric philosophy, mandate each agent to possess the non-trivial capability of network inference. In this manner, the more sophisticated network architectures of [23, 14, 24] inevitably introduce a heavier computational overhead.

On another front, we are aware of two recent endeavors [10, 25] that bear resemblance to our partially environment-centric approach. Ji *et al.* [10] introduced an augmented GNN policy, introducing environmental router nodes in addition to agent nodes. However, agents remain modeled as GNN nodes, necessitating agent-wise network inferences. On the other hand, Ye *et al.* [25] proposed a kernel-based neural policy wherein a small set of kernels dictates all agents' velocities. Nevertheless, this approach mandates agents to operate in a centralized manner, incurring substantial communication costs anew. In contrast, our method adheres solely to an environment-centric framework, demanding only rudimentary computational capabilities from the agents. Moreover, our method operates in a fully decentralized manner, eliminating the need for agent coordination and requiring them solely to consult the environment-centric GRNN policy for velocity modulation.

**Traffic-Rule Modeling:** The remarkable versatility, scalability, and efficiency exhibited by real-world traffic networks have spurred a substantial volume of research aimed at analyzing, emulating, and generalizing these networks to novel contexts. One avenue of inquiry revolves around data-driven traffic simulations and predictions [26, 27, 28]. Guided by real-world agent trajectory datasets, these methods train a probabilistic model to forecast agents' forthcoming trajectories. However, these probabilistic models remain oblivious to the inherent physical dynamics governing agents, often failing to ensure collision-free interactions among them. A complementary line of work leverages RL to enhance real-world traffic networks through traffic light control strategies [29].

Amidst these environment-centric approaches, traffic rules play an important role in agent-centric autonomous driving algorithms to ensure driving safety. For instance, Greenhalgh and Mirmehdi [30] introduced a model for text-based traffic sign recognition, while Franklin *et al.* [31] devised a framework to detect violations of traffic signals. Additionally, Lee *et al.* [32] formulated a method for predicting imminent freeway collisions. When compared with our method, it becomes evident that all the aforementioned techniques are heavily engineered toward the specific set of real-world traffic rules. In contrast, we advocate the utilization of a neural policy to encode an arbitrary set of unknown traffic rules, which can be learned to resolve congestion for arbitrary environmental and agent contexts. This dynamic approach not only engenders adaptability, but also accommodates scenarios where conventional, pre-defined rules may not suffice.

### III. PROBLEM STATEMENT

In this section, we present the formulation of the multi-agent navigation problem, addressing a scenario involving a swarm of agents operating in a 2D environment and we take the following standard assumption:

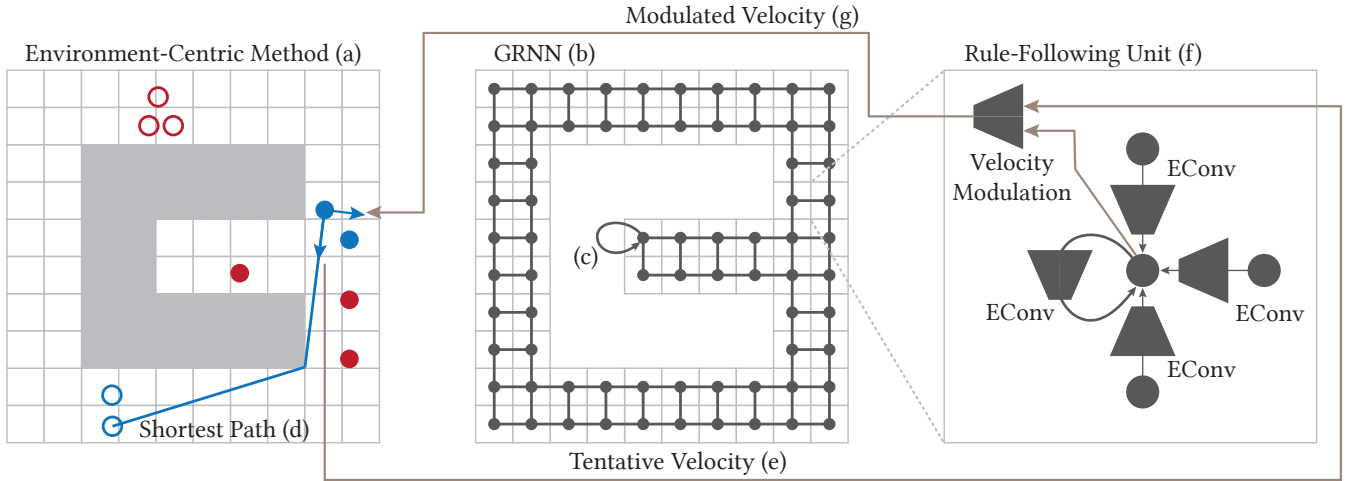


Fig. 1: Two distinct groups of agents (red and blue circles) are engaged in navigation, each with its designated goal positions (red or blue hollow circles). Our environment-centric technique (a) divides the environment into discrete blocks, upon which a GRNN (b) is deployed based on the block network. The GRNN inference process is executed just once, preceding the commencement of the navigation task, where the recurrent units update the internal state of each block and encode the pertinent traffic rules (c). During each timestep, an agent calculates the shortest path (d) to yield a tentative velocity (e). The agent then indexes the block it resides within, and proceeds to query the block’s rule-following unit (f), resulting in a congestion-resolving modulated velocity (g).

**Assumption 3.1:** All agents are circular-shaped with radius  $r$ . Agents are omnidirectional and capable of instantaneous change of moving direction. Agents move in a rectangular static environment with arbitrary interior obstacles, also of rectangular shapes.

The position of the  $i$ th agent at the  $t$ -th time step is denoted as  $x_i^t \in \mathbb{R}^2$ . The principal objective of each agent is to traverse from its initial position  $x_i^0$  to a predefined goal location  $g_i$ , while ensuring a trajectory devoid of collisions. Our approach is built off of the local navigation algorithms [7, 8], which offer a guarantee of collision-free motion. These algorithms take as input the desired velocity  $v_i^t$  for each agent at each time step, and then output the positions of the agents in the subsequent time step by approximating the solution to the following optimization problem:

$$\begin{aligned} \operatorname{argmin}_{x_i^{t+1}} \sum_i \|x_i^{t+1} - x_i^t - v_i^t\|^2 \\ \text{s.t. } \|x_i^{t+\gamma} - x_j^{t+\gamma}\| \geq 2r \quad \forall i \neq j \wedge \gamma \in [0, 1], \end{aligned} \quad (1)$$

where  $x_i^{t+\gamma}$  denotes the interpolated position of an agent at a fractional time instance  $t + \gamma$ . This optimization moves agents via the desired velocity as much as possible, while the associated constraints ensure a minimum separation distance of  $2r$  is maintained at every time instance.

Notably, local navigation algorithms are limited to ensuring collision avoidance, leaving the high-level behavior of the agents to be specified by the desired velocities  $v_i^t$ . The primary objective of our work lies in developing a neural navigation policy to predict  $v_i^t$  for all agents to facilitate navigation tasks while minimizing instances of congestion. In addition to the settings outlined above, our study incorporates two key assumptions, considering the constraints posed by existing swarm robot hardware [12, 13]. First, we operate under

the assumption that agents possess limited computational resources. Specifically, our requirement is solely for the agents to compute a local collision-free velocity, and no other substantial computations are permissible. Second, we consider the agents to possess restricted communication capabilities. In this context, agents are permitted to independently and asynchronously gather information from proximate computational nodes, while coordination of motion between any two agents is not allowed.

#### IV. NEURAL TRAFFIC RULE MODEL

Our approach is inspired by the inherent resemblance between real-world traffic networks and the specific challenge outlined in Section III. Real-world agents execute merely basic localized collision resolutions, steering clear of close-by pedestrians or the vehicle immediately ahead. Additionally, the entire set of agents navigates with minimal coordination. Our neural navigation policy is designed to mimic this paradigm. We optimize a parameterized set of traffic rules (Section IV-A) that are consistently followed by all agents in a decentralized manner (Section IV-B).

##### A. Neural Traffic Rule Parameterization

The comprehensive framework of our methodology is illustrated in Figure 1. We begin by recognizing that traffic rules are dictated solely by the local characteristics of the environment. For instance, the rules governing a crossroad remain unaffected by the type of road situated a block away. We incorporate this principle by segmenting the free space into square blocks. This arrangement constructs a block network, where each block maintains connections with its adjacent neighbors. To facilitate bidirectional traffic within each passage, we ensure that our block network has a minimum width

of 2 blocks (Figure 1b). Furthermore, a noteworthy feature is that traffic rules exhibit independence from individual agents. This attribute empowers a single set of rules to address arbitrary agent navigation scenarios without solving intricate decision-making problems. In view of this, we parameterize these agent-agnostic traffic rules via a GRNN as introduced by Ruiz *et al.* [33].

The core of our GRNN resides in the hidden-state representation, denoted as  $h_{ij}$ , for the  $ij$ th block. This representation is continually updated through an Edge Convolution (EConv) mechanism, iteratively realized as follows (Figure 1f):

$$h_{ij}^{k+1} = H(h_{ij}^k, \oplus_{kl \in \mathcal{N}_{ij}} \text{EConv}(x_{ij} - x_{kl}, h_{kl}^k)). \quad (2)$$

In this context,  $x_{ij}$  denotes the spatial coordinates of the center of the  $ij$ th block, and  $\mathcal{N}_{ij}$  represents its set of neighboring blocks. The function  $\text{EConv}(\bullet)$  is parameterized as a Multi-Layer Perceptron (MLP) that encodes the information from each neighboring block. The cumulative information propagation is subsequently summarized via the  $\oplus$  operator. We have experimented with  $\oplus$  being either concatenation or summation and the results are satisfactory in both cases. This aggregated information is then channeled through the hidden state update function  $H(\bullet)$ , itself parameterized by another MLP. Starting with  $h_{ij}^0 = 0$ , the GRNN progressively refines all  $h_{ij}^k$  until a fixed point  $h_{ij}^k \rightarrow h_{ij}^*$  is attained. The converged states implicitly encapsulate the essential traffic rules. In practical application, we approximate the convergence by performing  $K$  updates.

Since our GRNN is completely independent of agents, a single inference step is adequate to derive  $h_{ij}^K$  prior to the commencement of navigation tasks, ensuring runtime efficiency. Notably, various design approaches for our traffic rule network exist, and interested readers can refer to [34] for an overview of design insights. The main objective of our GRNN framework is to learn traffic rules from one environment and generalize them to different environments with distinct topologies in an agent-agnostic manner.

### B. Rule-Following Navigation Policy

With a well-defined set of traffic rules, agents can follow them to achieve navigation while avoiding congestion. The real-world agent navigation unfolds in two distinct stages. In the initial stage, such as a scenario involving a vehicle driver, a routing algorithm is engaged—examples include popular tools like Google Maps—to chart a course toward the desired destination. However, these routing algorithms primarily rely on a coarse estimate of travel time and do not deeply incorporate knowledge of traffic rules into their decision-making process. Subsequently, when the driver embarks on the selected route, the traffic rules take center stage. The driver’s journey is guided and informed by these regulations, ensuring smooth and unimpeded passage along the chosen route.

Our approach emulates this process through a two-stage navigation policy, denoted as  $v_i^t = \pi(x_i^t, g_i)$ . This policy takes the agent’s present location  $x_i^t$  and its designated goal position  $g_i$  as inputs and subsequently generates the rule-following, goal-directed velocity  $v_i^t$ .

During the first stage, our policy capitalizes on a visibility graph algorithm to ascertain the shortest trajectory connecting  $x_i^t$  and  $g_i$  (Figure 1d). The visibility graph precomputes all pairs of vertices that are visible to each other without being blocked by other parts of the environment. The computation of such graph has been well-studied and we use the efficient algorithm proposed in [35]. Subsequent to this precomputation, a single visibility query is conducted per agent, establishing connections between  $x_i^t$  and all visible graph nodes. A final shortest path query on the graph results in the sought-after trajectory. We denote this procedure as the shortest path function  $\bar{v}_i^t = \text{SP}(x_i^t, g_i)$ , giving the tentative velocity that propels the agent along the optimal path.

Following the initial stage, we utilize our GRNN to modulate  $\bar{v}_i^t$ , transforming it into the rule-compliant velocity  $v_i^t$ . Given that traffic rules are inherently localized, we find it sufficient to query the nearest block within which our agent resides, as is indexed by its center position:

$$I(x_i^t) = \lfloor x_i^t \rfloor + 0.5. \quad (3)$$

Subsequently, we enact the velocity modulation process via another MLP, define as the Rule-Following Unit (RFU):

$$v_i^t = \text{RFU}\left(h_{I(x_i^t)}^K, x_i^t - I(x_i^t), \bar{v}_i^t\right), \quad (4)$$

This MLP incorporates three crucial parameters: the hidden state  $h_{I(x_i^t)}^K$  of the underlying block, encoding the local traffic rule; the relative position between the agent and the block’s center  $x_i^t - I(x_i^t)$ ; and the tentative velocity  $\bar{v}_i^t$ . As an optional measure to ensure training stability and maintain runtime performance, we have the option to constrain the velocities  $\bar{v}_i^t$  and  $v_i^t$  using the maximum agent velocity  $v_{\max}$  via the clamping function:  $v = v / \max\{1, (\|v\| + \epsilon) / v_{\max}\}$  where  $\epsilon$  is a small regularization constant.

In real-world scenarios, agents are tasked with route selection, e.g., using Google Maps, followed by adherence to traffic rules. This paradigm underscores that the function  $\pi(\bullet)$  is executed on the agent part, with the environment only storing the  $h_{I(x_i^t)}^K$ . However, in the context of a robotic system, a more nuanced distribution of responsibilities becomes feasible. For agents with limited computational capabilities, the agent only needs to store  $x_i^t$  and  $g_i$ , and queries the environmental nodes for the desired  $v_i^t$ . For agents endowed with slightly more computational resources, agents can independently determine routes by running the  $\text{SP}(\bullet)$  function and command the environmental nodes to execute the  $\text{RFU}(\bullet)$  function. Significantly, regardless of the deployed setting, these agents operate in a fully decentralized manner, eliminating the need for inter-agent communication. Essentially, each agent is only required to identify the block and query the RFU.

### C. Policy Optimization

Harnessing the intrinsic locality of traffic rules, our policy is defined by three compact networks: H, EConv, and RFU. Nonetheless, optimizing this policy presents a formidable challenge due to the inherently non-smooth and stochastic nature of the navigation process. Due to the involvement of the non-differentiable block indexing operator (Equation 3),

conventional model-based differentiable policy optimization techniques such as [25] cannot be applied. Further, we consider intricate scenarios where agents can enter or exit the environment at arbitrary time points. Therefore, the collective state of agents undergoes dimension changes, deviating from the premises underpinning standard RL algorithms [36]. Finally, agents can undertake detours to follow traffic rules and avoid local collisions, making it difficult to assess task accomplishment and gauge the extent of congestion. These challenges are well-recognized in the realm of RL in the context of sparse rewards [37].

To confront these issues, we introduce two training approaches. When groundtruth traffic rules, denoted as  $\pi^*$ , are known to the user, we adopt the IL paradigm [38] to optimize our policy by emulating the expert's behavior. In this vein, we assemble a dataset of simulation scenarios, collectively denoted as  $\mathcal{S} = S_i$ . Each scenario,  $S_i$ , takes the form of a tuple:  $S_i = \langle F, S(\bullet), D(\bullet) \rangle$ . Here,  $F$  represents the environment's freespace geometry. The function  $\{x_i^t, g_i\} = S(x_{1,2,\dots}^t)$  generates new agents to enter the environment with respective goal positions, based on the current agent configuration. Similarly,  $\{x_i^t\} = D(x_{1,2,\dots}^t)$  determines which subset of agents have reached their goals and should exit the environment. During each iteration of IL training, we simulate navigations for a randomly selected scenario,  $S_i \in \mathcal{S}$  over a span of  $T$  timesteps. This simulation process accumulates a collection of agent transition tuples to form the dataset  $\mathcal{D} = \{ \langle x_i^t, g_i, \pi(x_i^t, g_i), \pi^*(x_i^t, g_i) \rangle \}$ . IL proceeds by minimizing the following expert discrepancy:

$$\mathcal{L}_{\text{IL}}(S_i, \theta) = \frac{1}{|\mathcal{D}|} \sum_{\langle x_i^t, g_i \rangle \in \mathcal{D}} |\pi(x_i^t, g_i) - \pi^*(x_i^t, g_i)|^2. \quad (5)$$

Here,  $x_i^t$  and  $g_i$  remain fixed variables, so the objective is differentiable under appropriate parameterization.

When an expert rule is not available, we train our policy using evolutionary RL [39]. Unlike conventional RL, this methodology bypasses the use of state-dependent rewards. Instead, it operates based on the assumption of a scenario-dependent sparse reward, denoted as  $R(S_i, \theta)$ , where  $\theta$  corresponds to our learnable parameters, and aims to maximize the expected reward:

$$\mathcal{L}_{\text{RL}}(S_i, \theta) = \mathbb{E}_{\epsilon \in \mathcal{N}(0, I)} [R(S_i, \theta + \epsilon)], \quad (6)$$

which essentially smoothens the reward signal. This technique iteratively updates  $\theta$  using a stochastic gradient estimator of the following form:

$$\begin{aligned} \nabla_{\theta} \mathcal{L}_{\text{RL}}(S_i, \theta) &= \mathbb{E}_{\epsilon \in \mathcal{N}(0, I)} [R(S_i, \theta + \epsilon) \epsilon] / \sigma \\ &\approx \frac{1}{n\sigma} \sum_{k=1}^n R(S_i, \theta + \epsilon_k) \epsilon_k, \end{aligned} \quad (7)$$

where  $\sigma$  denotes the estimated standard deviation and the estimator follows from log-likelihood technique and random sampling [40]. To enhance performance, we adopt fitness shaping as detailed in [40] and use the following alternative

approximation:

$$\begin{aligned} \nabla_{\theta} \mathcal{L}_{\text{RL}}(S_i, \theta) &\approx \frac{1}{\sigma} \sum_{k=1}^n u_k \epsilon_k, \\ u_k &\triangleq \frac{\max\left(0, \log\left(\frac{n}{2} + 1\right) - \log(k)\right)}{\sum_{j=1}^n \max\left(0, \log\left(\frac{n}{2} + 1\right) - \log(j)\right)} - \frac{1}{n}, \end{aligned} \quad (8)$$

where  $\{R(S_i, \theta + \epsilon_k)\}_{k=1}^n$  is the descending ordered sequence derived from  $\{R(S_i, \theta + \epsilon_k)\}_{k=1}^n$ . The fitness shaping is known to improve the convergence as recommended in [39]. Our overall algorithm follows a sequential routine, iteratively sampling of a batch of scenarios, executing simulations to assess  $R(S_i, \theta)$ , and forming the gradient estimator to update  $\theta$ .

The efficacy of the optimized policy hinges significantly on the reward signal  $R(S_i, \theta)$ . We discern that a well-constructed set of traffic rules should universally ensure freedom from congestion. That is, an agent can traverse from any starting point to any goal position without encountering congestion. Inspired by this insight, we tailor our reward to reflect the worst-case fraction of travel across all agents. In essence, we begin by identifying the set of agents that emerge during the simulation of  $S_i$ , denoted as  $\mathcal{A}(S_i) = \langle x_i, g_i \rangle$ . For each of these agents, we gauge their fraction of travel through the metric:  $F(x_i, g_i) = 1 - \text{SD}(x_i^T, g_i) / \text{SD}(x_i^0, g_i)$ , where  $\text{SD}(\bullet)$  signifies the shortest-distance function computed using the visibility graph, analogous to  $\text{SP}(\bullet)$ . Subsequently, we approximate the worst  $F(\bullet)$  over all agents, employing a soft-min function:

$$R(S_i, \theta) = \frac{\sum_{\langle x_i, g_i \rangle \in \mathcal{A}(S_i)} F(x_i, g_i) e^{-\alpha F(x_i, g_i)}}{\sum_{\langle x_i, g_i \rangle \in \mathcal{A}(S_i)} e^{-\alpha F(x_i, g_i)}}, \quad (9)$$

where  $\alpha$  is a softness-controlling parameter. To ensure training stability, we further incorporate curriculum learning that interleaves the increase of  $\alpha$  with the update of  $\theta$ . The complete workflow of our evolutionary RL is outlined in Algorithm 1.

---

### Algorithm 1 Evolutionary RL

---

**Input:** A set of scenarios  $\mathcal{S}$ , initial  $\theta, \alpha$ , parameter  $B, \eta$

**Output:** The universally congestion-free  $\pi(\bullet)$

- 1: **while**  $\theta$  not converged **do**
  - 2:    $\mathcal{D} \leftarrow \emptyset, \theta_0 \leftarrow \theta$ , sample  $S_i \in \mathcal{S}$     $\triangleright$  Batch size = 1
  - 3:   **for**  $i = 1, \dots, B$  **do**
  - 4:      $\epsilon \in \mathcal{N}(0, I)$ ,  $\theta \leftarrow \theta_0 + \epsilon$ ,  $\{x_i^0, g_i\} \leftarrow \emptyset$
  - 5:      $\triangleright$  Simulate scenario
  - 6:     **for**  $t = 0, \dots, T - 1$  **do**
  - 7:       Emit  $\{x_i^t, g_i\} \leftarrow \{x_i^t, g_i\} \cup S(x_{1,2,\dots}^t)$
  - 8:       Delete  $\{x_i^t, g_i\} \leftarrow \{x_i^t, g_i\} - D(x_{1,2,\dots}^t)$
  - 9:        $\triangleright$  Local navigation algorithm (LNav) [7, 8]
  - 10:       Simulate  $\{x_i^{t+1}, g_i\} \leftarrow \text{LNav}(\{x_i^t, g_i\})$
  - 11:     Collect  $\mathcal{A}(S_i)$  and  $\mathcal{D} \leftarrow \mathcal{D} \cup \{ \langle R(S_i, \theta), \epsilon \rangle \}$
  - 12:   Evaluate  $\nabla_{\theta} \mathcal{L}_{\text{RL}}(S_i, \theta)$  using  $\mathcal{D}$
  - 13:   Update  $\theta \leftarrow \theta_0 + \eta \nabla_{\theta} \mathcal{L}_{\text{RL}}$
  - 14:   Optionally increase  $\alpha$
  - 15: **Return**  $\theta$  and  $\pi(\bullet)$
-

## V. EVALUATION

Our method is realized through a robust implementation using C++ and PyTorch. We employ C++ to implement the implicit crowd algorithm [8], serving as our foundational local navigation framework. Our simulations run in parallel across multiple scenarios. For uniformity and consistency across all our experiments, we adopt a standardized network architecture. Our policy is parameterized through the integration of three compact MLPs. The function  $EConv(\bullet)$  is parameterized as:  $FC32 \rightarrow ReLU \rightarrow FC32 \rightarrow ReLU \rightarrow FC32$ . The function  $H(\bullet)$  is parameterized as:  $FC32 \rightarrow ReLU \rightarrow FC32 \rightarrow ReLU$  and we set  $|h_{ij}| = 32$ . Finally, the function  $RFU(\bullet)$  is parameterized as:  $FC32 \rightarrow ReLU \rightarrow FC32 \rightarrow ReLU \rightarrow FC16 \rightarrow ReLU \rightarrow FC2$ . For IL training, we set  $T = 500$  during each round of IL and run at least 1000 iterations of Adam to optimize  $\mathcal{L}_{IL}$  with a learning rate of  $10^{-4}$ . For RL training, we set  $T = 500$ ,  $\sigma = 0.02$ ,  $B = 10$ , the initial  $\alpha = 0$ , which is then reduced by 0.01 every 2000 iterations. Note that we set  $\eta = 2 \times 10^{-4}$  for the first 5000 iterations and  $\eta = 2 \times 10^{-5}$  for the rest. All experiments are performed on a single desktop machine with an AMD EPYC 7K62 CPU and we perform all computations on CPU, on which an IL training takes 48 hours and an RL training takes one week.

### A. Scenario Dataset

Our method requires two datasets: a training scenario dataset denoted as  $\mathcal{S}$  and a testing scenario dataset referred to as  $\mathcal{S}'$ , with  $\mathcal{S}'$  specifically comprising scenarios that have not been encountered during training. We set  $|\mathcal{S}|$  to be 85 and 120 for the training dataset of RL and IL, respectively; We set  $|\mathcal{S}'| = 30$  for the testing dataset of both IL and RL. To enable dynamic navigation within each scenario  $S_i$ , we further manually create a collection of starting and goal positions. During each time step, our seeding function  $S(\bullet)$  continually generates new agents, originating from these predefined starting positions as long as the newly generated agents do not collide with any existing ones. In addition, we set the maximum number of agents for each scenario to be 40 in the experiments. Conversely, our deleting function  $D(\bullet)$  removes agents that have successfully reached the same block as their respective goal positions reside.

### B. Performance

In our evaluation, we conduct a comparative analysis of our method against a baseline approach that relies solely on the shortest path without any rule-following modulation. This baseline policy, denoted as  $\pi^b$ , is achieved by setting  $v_i^t = \bar{v}_i^t$ . We gauge the performance of these policies using two fundamental metrics: the average fraction of travel  $R_0$  and the worst fraction of travel  $R_\infty$  defined as follows:

$$R_0 = \frac{1}{|\mathcal{S}'|} \sum_{S_i \in \mathcal{S}'} R_{\alpha=0}(S_i, \theta) \quad R_\infty = \frac{1}{|\mathcal{S}'|} \sum_{S_i \in \mathcal{S}'} R_{\alpha=\infty}(S_i, \theta).$$

Our initial evaluation focuses on the performance of IL-trained policies, which necessitates a set of groundtruth traffic rules. To this end, we craft a set of traffic rules emulating




Fig. 2: The set of allowed moving directions  $\mathcal{V}_{ij}$  for 4 types of blocks used by our expert policy. We do not label the  $\mathcal{V}_{ij}$  for boundary blocks, which depend on the boundary conditions therein.

real-world road networks. Each lane is designed to accommodate travel in only one specific direction. Furthermore, we incorporate the conventions of real-world roundabouts at road intersections, wherein all vehicles are constrained to move in a counter-clockwise direction. This set of rules yields a predefined collection of allowable moving directions, denoted as  $\mathcal{V}_{ij}$  for the  $ij$ -th block. A visual representation of these directions is provided in Figure 2. Our expert policy then determines the modulated velocity as one of the allowed moving directions that aligns best with  $\bar{v}_i^t$  as defined through the following expression:

$$\pi^*(x_i^t, g_i) = \operatorname{argmax}_{v \in \mathcal{V}_I(x_i^t)} \langle v, \bar{v}_i^t \rangle.$$

To track the convergence of our IL training process, we present a convergence history plot in Figure 3. Notably, this plot illustrates that IL rapidly converges often requiring as few as 50 rounds. Following training, we summarize the performance of all three policies in Table I. It is evident that there exists a minor performance gap in terms of  $R_0$  between  $\pi^{IL}$  and  $\pi^*$ , which provides confirmation that our meticulously parameterized neural policy possesses the capability to effectively capture and express a subset of real-world traffic rules. However,  $\pi^{IL}$  is vulnerable to variations in  $R_\infty$  when evaluated in new scenarios. This limitation is mitigated by  $\pi^{RL}$ , which does not strictly adhere to  $\pi^*$  and thus has greater flexibility in adapting to unfamiliar situations. Furthermore, both  $\pi^{IL}$  and  $\pi^*$  exhibit notably superior performance compared to  $\pi^b$ , demonstrating the remarkable effectiveness of incorporating real-world traffic rules in alleviating congestion-related issues.

TABLE I: Testing score of baseline  $\pi^b$ , RL-trained  $\pi^{RL}$ , IL-trained  $\pi^{IL}$ , and the expert  $\pi^*$  in terms of  $R_0$  and  $R_\infty$ . We run all the test scenarios 10 times and report the average  $\pm$  deviation.

	$\pi^b$	$\pi^{IL}$	$\pi^{RL}$	$\pi^*$
$R_0$	$25.35 \pm 3.35$	$33.97 \pm 1.98$	$38.36 \pm 0.52$	$38.45 \pm 0.07$
$R_\infty$	$5.13 \pm 2.33$	$9.31 \pm 6.75$	$37.78 \pm 0.18$	$38.06 \pm 0.10$

Moving on, our evaluation proceeds to the RL-trained policy. Again we plot the convergence history in Figure 4 and summarize its performance in the second column of Table I.




Fig. 3: The convergence history in terms of  $R_0$  and  $R_\infty$  plotted against the index of IL round. We show the shadow plot of 3 runs to highlight the consistency of our method.




Fig. 4: The convergence history in terms of  $R_0$  and  $R_\infty$  plotted against the index of RL iteration. We show the shadow plot of 3 runs to highlight the consistency of our method.

Some instability is observed during the training process as we use smaller values of  $\alpha$  leading to progressively non-smooth reward signals. In addition, this behavior can be attributed to the small value of  $B = 10$  causing imprecise approximation of the policy gradient. Conversely, opting for larger values of  $B$  can significantly escalate the training cost. Upon completion of training, we observe a distinguished improvement in terms of  $R_\infty$  compared to  $\pi^{IL}$ . This result is consistent with our expectations, as RL autonomously searches for optimal traffic rules without requiring expert guidance, thereby enhancing robustness. As anticipated, both the IL and RL-trained policies substantially outperform  $\pi^b$ . This underscores not only the expressiveness of our method but also its capacity to autonomously discover suitable traffic rules, guided solely by congestion-resolving reward signals. A visual comparison is given in Figure 5.

Ultimately, we evaluate the runtime inference costs of our approach and the agent-centric method [10] by comparing




Fig. 5: We show congested scenarios caused by using the baseline  $\pi^b$  (left) and the corresponding congestion-free scenarios achieved by our  $\pi^{RL}$  (right). Agents are colored according to their group.




Fig. 6: The runtime total inference cost profiled in milliseconds, which is averaged over the 10 runs on test scenarios.

them in Figure 6. Both methodologies utilize GNNs to parameterize the policy. For this comparison, we pick a random instance from our scenario with a variable number of agents. We then run the GNN inference of the agent-centric method on the agent network, using the same parameter setting as [10]. Remarkably, our architecture conducts GRNN inference only once, leading to significantly enhanced performance. As depicted in Figure 6, we attain a peak performance of 400 milliseconds while managing a cohort of 240 agents.

## VI. CONCLUSION & FUTURE WORKS

We introduce an environment-centric architecture for neural navigation policies, designed to replicate real-world traffic networks. Our approach demonstrates the feasibility of training such a policy using IL guided by expert traffic rules or RL with congestion-resolving reward signals. Our approach eliminates the need for runtime inter-agent communication, resulting in significantly reduced computational burden compared to agent-centric navigation policies.

This research opens up a plethora of promising avenues for future exploration, such as hardware deployment. One notable aspect to consider is the modeling of time-dependent traffic rules, such as those associated with traffic lights. Furthermore, our current method exclusively accounts for homogeneous traffic rules, yet in the real world, rule variations are contingent upon the context of the environment, encompassing scenarios like parking lots and highway checkpoints. Extending our

approach to encompass these more nuanced factors represents an intriguing area for future research endeavors. Finally, our training procedure is time-consuming and CPU-bound. This is attributed to our collision-free agent simulator Equation 1, which involves solving a joint optimization problem encompassing all agents. We plan to utilize parallel computations to speed up this procedure and expedite the training process.

## REFERENCES

- [1] K. Azadeh, R. De Koster, and D. Roy, "Robotized and automated warehouse systems: Review and recent developments," *Transportation Science*, vol. 53, no. 4, pp. 917–945, 2019.
- [2] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A survey of autonomous driving: Common practices and emerging technologies," *IEEE access*, vol. 8, pp. 58 443–58 469, 2020.
- [3] B. N. Silva, M. Khan, and K. Han, "Towards sustainable smart cities: A review of trends, architectures, components, and open challenges in smart cities," *Sustainable cities and society*, vol. 38, pp. 697–713, 2018.
- [4] B. De Wilde, A. W. Ter Mors, and C. Witteveen, "Push and rotate: Cooperative multi-agent path planning," in *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, 2013, pp. 87–94.
- [5] J. Yu and S. LaValle, "Structure and intractability of optimal multi-robot path planning on graphs," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27, 2013, pp. 1443–1449.
- [6] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [7] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *2008 IEEE international conference on robotics and automation*, Ieee, 2008, pp. 1928–1935.
- [8] I. Karamouzas, N. Sohre, R. Narain, and S. J. Guy, "Implicit crowds: Optimization integrator for robust crowd simulation," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–13, 2017.
- [9] T. Fan, P. Long, W. Liu, and J. Pan, "Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 856–892, 2020.
- [10] X. Ji, H. Li, Z. Pan, X. Gao, and C. Tu, "Decentralized, unlabeled multi-agent navigation in obstacle-rich environments using graph neural networks," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2021, pp. 8936–8943.
- [11] A. Puente-Castro, D. Rivero, A. Pazos, and E. Fernandez-Blanco, "A review of artificial intelligence applied to path planning in uav swarms," *Neural Computing and Applications*, pp. 1–18, 2022.
- [12] H. Wei, Y. Cai, H. Li, D. Li, and T. Wang, "Sambot: A self-assembly modular robot for swarm robot," in *2010 IEEE International Conference on Robotics and Automation*, IEEE, 2010, pp. 66–71.
- [13] M. Rubenstein, C. Ahler, and R. Nagpal, "Kilobot: A low cost scalable robot system for collective behaviors," in *2012 IEEE international conference on robotics and automation*, IEEE, 2012, pp. 3293–3298.
- [14] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 11 785–11 792.
- [15] J. Yu, "Average case constant factor time and distance optimal multi-robot path planning in well-connected environments," *Autonomous Robots*, vol. 44, no. 3-4, pp. 469–483, 2020.
- [16] K. Solovey, J. Yu, O. Zamir, and D. Halperin, "Motion planning for unlabeled discs with optimality guarantees," *arXiv preprint arXiv:1504.05218*, 2015.
- [17] S. Yang, T. Li, X. Gong, B. Peng, and J. Hu, "A review on crowd simulation and modeling," *Graphical Models*, vol. 111, p. 101 081, 2020.
- [18] L. He, J. Pan, S. Narang, and D. Manocha, "Dynamic group behaviors for interactive crowd simulation," in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2016, pp. 139–147.
- [19] L. He, Z. Pan, and D. Manocha, "Real-time decentralized navigation of nonholonomic agents using shifted yielding areas," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3656–3662.
- [20] M. Schuerman, S. Singh, M. Kapadia, and P. Faloutsos, "Situation agents: Agent-based externalized steering logic," *Computer Animation and Virtual Worlds*, vol. 21, no. 3-4, pp. 267–276, 2010.
- [21] T. Huang, S. Koenig, and B. Dilkina, "Learning to resolve conflicts for multi-agent path finding with conflict-based search," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, 2021, pp. 11 246–11 253.
- [22] S. D. Han and J. Yu, "DDM: Fast near-optimal multi-robot path planning using diversified-path and optimal sub-problem solution database heuristics," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1350–1357, 2020, Note: presented at ICRA 2020.
- [23] Q. Tan, T. Fan, J. Pan, and D. Manocha, "Deepmnavigate: Deep reinforced multi-robot navigation unifying local & global collision avoidance," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 6952–6959.
- [24] J. Blumenkamp and A. Prorok, "The emergence of adversarial communication in multi-agent reinforcement learning," in *Conference on Robot Learning*, PMLR, 2021, pp. 1394–1414.
- [25] X. Ye, Z. Pan, X. Gao, K. Wu, and B. Ren, "Differentiable learning of scalable multi-agent navigation policies," *IEEE Robotics and Automation Letters*, vol. 8, no. 4, pp. 2229–2236, 2023.
- [26] Q. Chao, Z. Deng, J. Ren, Q. Ye, and X. Jin, "Realistic data-driven traffic flow animation using texture synthesis," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 2, pp. 1167–1178, 2017.
- [27] Y. Ma, X. Zhu, S. Zhang, R. Yang, W. Wang, and D. Manocha, "Trafficpredict: Trajectory prediction for heterogeneous traffic-agents," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, 2019, pp. 6120–6127.
- [28] S. Suo, S. Regalado, S. Casas, and R. Urtasun, "TrafficSim: Learning to simulate realistic multi-agent behaviors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10 400–10 409.
- [29] K.-L. A. Yau, J. Qadir, H. L. Khoo, M. H. Ling, and P. Komisarczuk, "A survey on reinforcement learning models and algorithms for traffic signal control," *ACM Computing Surveys (CSUR)*, vol. 50, no. 3, pp. 1–38, 2017.
- [30] J. Greenhalgh and M. Mirmehdi, "Recognizing text-based traffic signs," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 3, pp. 1360–1369, 2014.
- [31] R. J. Franklin *et al.*, "Traffic signal violation detection using artificial intelligence and deep learning," in *2020 5th international conference on communication and electronics systems (ICCES)*, IEEE, 2020, pp. 839–844.
- [32] C. Lee, B. Hellenga, and F. Saccomanno, "Real-time crash prediction model for application to crash prevention in freeway traffic," *Transportation Research Record*, vol. 1840, no. 1, pp. 67–77, 2003.
- [33] L. Ruiz, F. Gama, and A. Ribeiro, "Gated graph recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 68, pp. 6303–6318, 2020.
- [34] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," *Advances in neural information processing systems*, vol. 30, 2017.
- [35] T. Lozano-Pérez and M. A. Wesley, "An algorithm for planning collision-free paths among polyhedral obstacles," *Communications of the ACM*, vol. 22, no. 10, pp. 560–570, 1979.
- [36] E. Vinitzky *et al.*, "Benchmarks for reinforcement learning in mixed-autonomy traffic," in *Conference on robot learning*, PMLR, 2018, pp. 399–409.
- [37] J. Hare, "Dealing with sparse rewards in reinforcement learning," *arXiv preprint arXiv:1910.09281*, 2019.
- [38] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
- [39] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, "Evolution strategies as a scalable alternative to reinforcement learning," *arXiv preprint arXiv:1703.03864*, 2017.
- [40] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber, "Natural evolution strategies," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 949–980, 2014.