

CAV Unsignalised Intersection Crossing with Unreliable Emergent Communication

Anastasios N. Kontogiorgis¹, Melanie Bouroche¹

¹*School of Computer Science and Statistics, Trinity College Dublin, Ireland*

Abstract

Reliable communication is crucial for effective and safe coordination among connected autonomous vehicles (CAVs), especially in complex scenarios such as roundabouts or intersections. This work investigates the effect of unreliable communication on emergent communication (EC) for coordinating autonomous vehicles at non-signalised intersections. Existing EC solutions typically assume reliable – error and noise free – communication, which is usually not the case in realistic scenarios. We evaluate how communication limitations such as message noise, partial and whole message loss affect the performance of four state-of-the-art models, namely CommNet, TarMac, IC3Net and GA-Comm in a non-signalised intersection task of increasing difficulty and reduced visibility. We investigate each model’s resilience to these communication disturbances, additionally analysing the comparative impact of each disturbance type and intensity on model success rates.

Keywords

Emergent communication, multi-agent, non-signalised intersection, unreliable communication, communication noise, message drops

1. Introduction

Intelligent transportation systems (ITS) aim to enable safer and more sustainable transportation by alleviating current mobility issues such as accidents, pollution and inefficient utilisation of resources [1]. Among the ITS technologies, vehicle-to-vehicle (V2V) communication has the potential to revolutionise the transportation sector and enable improvements in safety, energy efficiency, and infrastructure utilisation [2, 3, 4]. More specifically, information sharing among vehicles through vehicular ad hoc networks (VANETs) in which Connected Autonomous Vehicles (CAVs) communicate through V2V communication, and road-side infrastructure through vehicle-to-infrastructure communication (V2I), can extend situational awareness and assist in building a richer representation of the CAVs’ extended neighbourhood [5].

An especially challenging subset of traffic scenarios arises when vehicles need to coordinate in order to use a common resource, such as a roundabout or an intersection. Intersections are among the most complex and inefficient elements of current traffic systems [6] in which a disproportionate number of accidents, injuries and fatalities occur [7, 8]. Traditional solutions in autonomous driving manually define which actions to perform in specific situations using techniques such as behaviour-trees or finite state machines [9, 10]. While successful, these approaches lack the ability to generalise and cater for unexpected situations [11]. Designing ‘universal’ hand-crafted rules that can handle the complexity of all possible scenarios accounting for uncertainty and the intricate relations that emerge between road actors becomes a daunting task. Indeed, specifying a-priori intelligent behaviour in complex systems is considered challenging, if not impossible [12, 13].

Multi-agent reinforcement learning (MARL) provides a powerful framework for tackling problems in which the joint actions of multiple decision-making agents influence a shared environment. By modelling the interactions between agents, MARL has the potential to develop autonomous systems capable of navigating complex environments and collaborating to solve challenging tasks [14]. In

ATT’24: Workshop Agents in Traffic and Transportation, October 19, 2024, Santiago de Compostela, Spain

✉ kontogia@tcd.ie (A. N. Kontogiorgis); melanie.bouroche@tcd.ie (M. Bouroche)

🌐 <https://github.com/AnastasiosKo/NoisyComm> (A. N. Kontogiorgis); <https://www.scss.tcd.ie/Melanie.Bouroche/> (M. Bouroche)

🆔 0000-0003-4921-3835 (A. N. Kontogiorgis); 0000-0002-5039-0815 (M. Bouroche)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

particular, MARL has been proposed as a promising solution for coordinating connected autonomous vehicles (CAVs) in complex traffic scenarios [15, 16]. A key advantage of MARL is its ability to learn from experience and generalise to unexpected situations, rather than relying on predetermined rules. This flexibility allows for the development of strategies that can better handle complex and dynamic traffic scenarios.

Recently, *emergent communication (EC)* between reinforcement learning (RL) agents has gathered significant interest in the research community since the pioneering works of [17, 18]. In this approach, agents learn to coordinate through a shared channel, allowing for the discovery of communication protocols based on task requirements. Learned communication tends to be more flexible and goal oriented leading to improvements in coordination and task success [19, 20]. While significant advances have been made in the field, showing promise for solving real-world problems, achieving robust performance requires addressing communication reliability. In realistic environments the quality of communication is subject to changes due to interference such as noise, message jumbling, information congestion, message delays and losses [21]. While real-world constraints, namely limited bandwidth, communication bottleneck and efficiency issues, have been addressed in the literature [22, 23, 24, 25], communication itself is largely assumed to be error and noise-free [26, 27].

To address this limitation, this work investigates the behaviour of existing EC solutions - typically assuming reliable communication - in the presence of noise and message drops. We aim to address this gap by methodically testing the following state-of-the-art EC solutions, CommNet [18], IC3Net [24], TARMAC [22] and GA-Comm [28], analysing how communication constraints affect task performance in a non-signalised intersection environment. We investigate model robustness in the presence of noise and the impact of each type of communication disturbance on task performance.

This work makes the following contributions:

- A review and critique of existing literature in emergent communication tackling noise.
- A systematic investigation of the effects of unreliable communication on agent performance in the widely used non-signalised intersection environment first introduced by [18].
- A comparative analysis of model resilience to communication disturbances - disturbances not present during training - detailing the comparative impact of each disturbance type and intensity on model performance.

The rest of this paper is organised as follows. Section 2 provides a review of emergent communication applied in non-signalised intersections, and of the current literature on emergent communication in the presence of noise. Section 3 details the experimental setup including the non-signalised intersection environment used for evaluation and the noise models that emulate communication disturbances. In Section 4, the findings of training and testing with noise and message drops are presented and analysed, along with a discussion on the comparative impact of noise type and intensity on the performance of the models tested. Finally, Section 5 summarises the key findings and concludes the paper with a discussion of future work.

2. Background

Emergent communication among reinforcement learning agents is an evolving field that aims to develop communication protocols without prior explicit instructions or predefined language rules, allowing agents to collaborate towards a common goal. This approach enables the development of communication strategies based on the requirements of the task, providing increased flexibility. A question that naturally arises is whether (and how well) agents can learn a 'language' over a joint communication channel, allowing them to maximise their utility [29].

A large body of work in the field has been directed toward 'language'-based coordination of deep RL agents in complex tasks. Early works [30, 31] utilised predefined communication protocols to facilitate information exchange and collaboration, however, such approaches may prove rigid. In recent years learnable communication has been widely explored with the emerging communication protocols




Figure 1: Easy version of the traffic junction environment, agents need to advance through the two intersecting one-way lanes without collisions using actions 'accelerate' or 'brake', visibility is set to one grid cell and the maximum number of agents at any given time is set to 5. (Image adapted from [24]).

collaboratively solving tasks such as riddles [17], navigation in complex 3D environments [22] and agent coordination in non-signalised intersection environments [18, 24, 22, 28, 23].

The following sections explore the use of emergent communication in non-signalised intersection, highlighting state-of-the-art approaches that enhance agent coordination and communication efficiency. We further discuss the impact of communication constraints, and the concept of noise in emergent communication, broadly categorising relevant literature based on the type of noise it addresses.

2.1. Non-signalised Intersection Environment

For training and validation, we employ the non-signalised intersection first introduced by [18], a widely-used [24, 22, 28, 23] environment for bench-marking emergent communication models. It comprises intersecting pathways and agents with default vision of a 3×3 surrounding grid, limiting the visual range affects the agent ability to navigate the intersection necessitating communication to avoid collisions. Agents enter the intersection with probability p_{arrive} with the maximum number of cars at any moment given by N_{max} , which changes according to the level of difficulty. Each agent occupies a single grid cell per time-step and has an available action space of 'accelerate' (advancing one cell) or 'brake' (no move). The reward consists of a penalty -0.01τ that accumulates linearly over time and a collision penalty $r_{\text{collision}} = -10$ which classifies the episode as a failure. Success rate, used as an evaluation metric, indicates whether a collision occurred during an episode. The total reward at time t is given by:

$$r(t) = C_t r_{\text{collision}} + \sum_{i=1}^{N^t} \tau_i r_{\text{time}},$$

where C_t is the number of collisions at time t , and N^t is the number of cars present.

We focus on the easy level, shown in Figure 1, which features a pair of two one-way lanes within a 7×7 grid, accommodating a maximum of five vehicles ($N_{\text{max}} = 5$, $p_{\text{arrive}} = 0.3$).

2.2. Emergent Communication for non-signalised Intersection Crossing

Communication is an important aspect in intersection environments where agents need to coordinate their actions to avoid collisions and optimise traffic flow. Various techniques have been proposed to enable agents to learn when, what, and to whom to communicate, enhancing cooperation, communication efficiency and task performance.

One of the pioneering works in the field, CommNet [18] introduced a multi-pass communication framework for exchanging continuous-valued messages containing averaged transmissions of encoded

hidden states, either by broadcasting or to agents within a certain range; communication in the latter case can be represented as a dynamic graph. Testing in the non-signalised intersection with agent visibility set to zero, CommNet achieved a 90% average success rate of no collisions within an episode indicating that agents successfully use the emerging communication to coordinate. In IC3Net [24], a gating mechanism to control the communication action (i.e. whether the agent will communicate or not at the next step) and individualised rewards are additionally introduced, extending applicability to non-cooperative scenarios and improving scalability. Focusing on selective communication, TARMAC [22] uses attention to determine the recipients of goal-specific messages and enables dynamic team sizes inside which agents can communicate. Leveraging the representation power of graphs, [28] models agent relationships using a complete graph and two-stage – hard and soft – attention to detect and assess the importance of interactions between agents, allowing the model to dynamically adapt to the complexities of the environment by focusing on relevant interactions. The authors further extend this approach to a communication model (GA-COMM) by allowing each agent to attend to the messages of others when making decisions. Similarly, MAGIC [23] employs graph-attention to target communication and for message processing. Evaluating in the non-signalised intersection showed high success rates in difficult settings with reduced visibility for both approaches.

These approaches assume reliable communication, which is not realistic for practical scenarios. Messages can be limited in size, due to bandwidth limitations, and range. Additionally, unreliable connections can introduce various forms of interference such as noise corrupting message content, delays, message jumbling -messages reach the agents mixed-up in content- and losses, disrupting information exchange and accuracy. These constraints can impact the learning performance of agents and task success, presenting a challenge that needs to be addressed. Several studies have proposed techniques that optimise the communication process by either addressing specific agents [22] or utilising Networked MARL (NMARL) [32] allowing communication with neighbouring agents only. Other studies focus on reducing communication overhead by enabling agents to choose whether to communicate or not [24, 33, 34]. While such work addresses communication quality from the perspective of efficiency, it similarly assumes limited-capacity yet reliable error-free communication and does not consider underlying communication channel characteristics such as noise.

2.3. Learning Communication in the Presence of Noise

When considering noise in emergent communication, studies can be broadly categorised into two groups: (1) those leveraging communication to minimise the impact of noise in agent observations: for example by sharing information, agents can reach an agreement on the state of a complex environment [35]; and (2) those that address communication channel noise which affects the reliability of exchanged messages. We will focus on the latter category.

This group of studies relates to Levels A and C of Shannon and Weaver[36], referred to as the 'technical' and 'effectiveness' problems, focusing on how agents can achieve coordination by learning to communicate over a noisy channel. In particular, [37] address a cooperative task involving two agents communicating over a noisy link. The authors employ a joint strategy for simultaneously learning both communication and action selection, leading to improved performance and resulting in a learnt communication scheme that incorporates both data compression and error protection. The coordination task in this setup, however, does not explicitly depend on communication, as agents can independently be trained to navigate to the goal -which is fixed to a specific position across all episodes. Similarly, [38] explore the problem of a guide coordinating a scout over a noisy communication link. Here, the optimal policy is learnt by taking into account channel limitations instead of assuming perfect communication and subsequently employing a communication protocol to convey the actions of the guide. While this approach demonstrates effective learning under noise, its applicability to larger environments involving more complex tasks is not described by the authors. Scaling to a larger number of agents would require innovations in the structure of the learnt communication model [39].

The literature so far has shown that agents can effectively coordinate in the presence of noisy communication. Notably, the 'language' that emerges in such conditions is distinct from that which

arises in scenarios with reliable communication [26] and outperforms conventional communication [40]. Nonetheless, these studies typically focused on relatively simple tasks involving two agents and a single form of noise. A question that arises is how these adaptive communication strategies scale to complex scenarios of increasing difficulty with multiple agents coordinating while communication is affected by a range of interferences.

In the next part of this paper, we investigate how emergent communication models applied in a non-signalised intersection environment perform in the presence of unreliable communication.

3. Methodology

For a comprehensive evaluation of how different emergent communication mechanisms cope with unreliable communication, we compare state-of-the-art EC models CommNet, TarMac, IC3Net and GA-Comm in multi-agent cooperative scenarios on the commonly used traffic junction environment [18], by introducing noise and random message drops to address the following questions:

- What are the effects of unreliable communication on task performance for the above-mentioned models in a cooperative intersection scenario? Essentially, how do models cope in the presence of noise?
- What is the impact of each type of communication disturbance on model performance?

We further discuss the design and experimental setup for evaluating in the presence of unreliable communication, this includes detailing the training and testing settings, the noise models employed to simulate communication disturbances, and evaluation metrics used to assess model performance.

3.1. Design

To maintain consistency we have adopted parameter values and a training method that align with the original studies. Models are trained using multi-threaded synchronous policy gradient [24], each thread runs batch learning with a batch size of 500 and performs 10 weight updates per epoch. The RMSProp optimizer is employed with a learning rate α of 1×10^{-3} , a discount factor γ of 1.0, and a value coefficient β of 1×10^{-2} . More specifically, an additional value head is used in the policy network to estimate a value function, $V_\phi(o_t^i)$, at each agent’s observation o_t^i . Along with optimising the discounted total rewards, the training process also minimises the squared error of the estimated value, which acts a baseline, against the Monte Carlo predicted value. This process is balanced by coefficient β . The overall loss function $L(\cdot)$, and the policy function, $\pi_\theta(a_t^i|o_t^i)$, share most of the parameters θ and ϕ , except those in the policy and value head. Each agent’s LSTM hidden state size is set to 128 dimensions and subsequently, the message size is 128 dimensions. We use two rounds of communication as empirically this has shown to provide better performance and training speed in reliable communication conditions [41]. Finally, we run the traffic junction experiment for 2000 epochs and for a single seed. Post-training we assess how each trained model generalises in the presence of unreliable communication for 1000 epochs.

We use success rate as an evaluation metric signifying no collisions within an episode and limit agent vision to size 1 to increase difficulty and promote communication among the agents. Each experiment employs curriculum learning to facilitate training. More specifically, p_{arrive} is maintained at the initial value for the first 250 epochs, after which it is gradually increased between epochs 250 and 1250 to its final value 0.3, where it is maintained until the end, as shown in Table 1

For validation, we employ 1000 epochs under a difficulty setting analogous to training, featuring two one-way lanes and an identical vision range. The difficulty increase through p_{arrive} is proportionally aligned with the training phase to focus on model’s ability to cope in the presence of communication noise. Specifically, we adjust p_{arrive} for the initial 125 epochs as easy, the next 500 epochs as harder, and the final 375 epochs as the hardest. This scaling preserves the ratio of difficulty progression during training, with the Easy phase representing 12.5% of the total epochs, the Harder phase 50%, and

Table 1
Training and testing settings for the traffic junction task

Mode	Epochs	Difficulty	Add rate min	Add rate max	N-total	Current start	Current end
Train	2000	easy	0.1	0.3	5	250	1250
Test	1000	easy	0.1	0.3	5	125	625

the Hardest phase 37.5%. These settings are designed to isolate the impact of communication noise on performance by maintaining a consistent difficulty increase, therefore any observed performance degradation is attributed to the introduction of noise rather than the difficulty levels.

3.2. Noise Models

As discussed in Section 2.2 messages can become corrupted by noise, lost, jumbled in content or delayed. We address the first two cases, where noise is introduced into the communication tensor, altering message values, and where entire messages or parts of them are lost. Noise is applied to the communication tensor, and the effects on model success rates are measured against noiseless communication. The specific parameters and probabilities associated with noise and message drops are detailed in Table 2, which provides an overview of the different noise models and their impact on communication.

- **Gaussian noise** is added to the signal with a mean of μ and a controlled intensity level represented by the standard deviation σ , denoted as $N(\mu = 0, \sigma)$. Noise intensity is regulated by adjusting σ , which affects the spread or 'width' of the noise injected into the signal. This type of noise is used to simulate noise that comes from natural sources, such as thermal noise or interference.
- **Uniform Noise**, also known as random valued noise, distributes noise uniformly across a given range creating a uniform distribution of noise values. The noise distribution is generated within the range $[a, b]$, where a is low and b is high. This noise model is used to simulate random events, such as random bit errors in a communication system.
- **Partial Message Drop** simulates the scenario where parts of the messages are lost during transmission with probability P_{partial} . A mask is generated using a Bernoulli distribution with a probability equal to P_{partial} for each element of the communication tensor, c . The mask is then applied to the communication tensor, resulting in a partial loss of message parts.
- **Whole Message Drop** simulates the condition where all messages to one or more agents are dropped entirely with probability P_{whole} . A mask is created with probabilities equal to P_{whole} for each agent and applied to the communication tensor c . This results in the potential loss of whole messages to an agent or agents.
- **Combined Partial and Total Message Loss** simulates the scenario where both partial and total message drops may occur, denoted by P_{both} which is the product of P_{whole} and P_{partial} :

$$P_{\text{both}} = P_{\text{whole}} \times P_{\text{partial}}$$

We test with a wide range of noise intensities and message loss frequencies, as shown in Table 2, to evaluate and compare the impact of each noise type on task performance and model robustness. Disturbances are tested independently, meaning that Gaussian, Uniform noise and message loss are not introduced simultaneously.

4. Results and Analysis

In this section, we first evaluate the training performance of CommNet [18], TarMac [22], IC3Net [24], and GA-Comm [28] in the traffic junction task under ideal communication conditions. Subsequently,

Table 2


Summary of noise parameters and the associated probabilities for partial or total message loss used in the experiments.

Gaussian Noise	Uniform Noise	Partial Message Loss	Whole Message Loss	Combined Loss
$\sigma = 0.2, 0.5, 0.8$	$U[0.0, 0.5], U[0.0, 0.8]$	$P_{\text{partial}} = 0.4, 0.7$	$P_{\text{whole}} = 0.3, 0.6$	$P_{\text{both}} = 0.12, 0.42$

model performance is bench-marked in scenarios with and without the presence of unreliable communication. We investigate the models’ ability to overcome noise and identify which models are more robust to various communication disturbances. Finally, we analyse the effect of different noise types on task success, particularly focusing on the comparative impact of each noise type on model performance.

4.1. Training and Validation

Training and testing on the traffic junction task at the first level of difficulty and with a maximum of five agents produced high average success rates for most of the models. Notably, TARMAC shows some performance degradation after epoch 250, the point at which difficulty progressively increases. Figure 2a (a) illustrates the success rates across the first 1000 epochs of training.



(a) Success rates during training for the traffic junction task. (b) Success rates during testing assuming reliable communication.

Figure 2: Comparison of success rates during training and testing.

High average success rates were also maintained throughout the testing phase including during intervals of increased difficulty. From epochs 125 to 625 there is a gradual increase in vehicle add rates, and from epoch 625 until the end, vehicle add rates reach their maximum. The narrow range between minimum and maximum success rates throughout all models during testing indicates consistent performance across epochs. CommNet and GA-Comm showed the best performance in both training and testing. In contrast, TARMAC exhibits the lowest average success rate in training and the largest range between minimum and maximum in both training and testing, see Figure 2b (b). This variation suggests that TARMAC might be more sensitive to changes in p_{arrive} or that more epochs are required to stabilise its performance, which will be explored in future work. The success rate values obtained from testing the trained models with reliable communication will serve as a benchmark to quantify the impact of communication disturbances. This will help identify each model’s ability to generalise in the presence of unreliable communication.


4.2. Unreliable Communication

Following the evaluation of model performance in ideal communication conditions, we introduce noise and message drops of various intensities (as detailed in Table 2) into the communication tensor. These values, designed to simulate low to high levels of interference, were chosen to provide a balance between realistic noise levels and challenging the model’s ability to adapt. Beginning with CommNet, which previously demonstrated the highest average success rate in testing without noise, we identify the

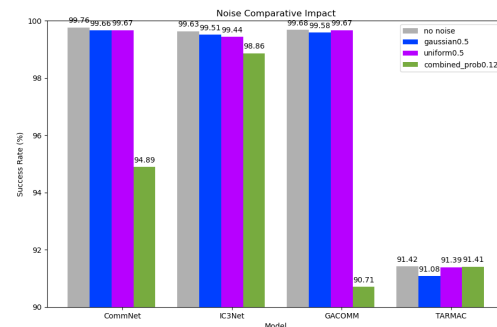
Table 3

Success rates and performance degradation for CommNet across all noise types and intensities.

Noise Type	Intensity	Average Success (%)	Degradation (%)
No noise	-	99.7664	-
Gaussian	0.3	99.72	0.04
	0.5	99.66	0.1
	0.8	99.43	0.33
Uniform	0.5	99.67	0.09
	0.8	99.52	0.23
Whole	0.3	95.19	4.57
	0.6	92.1	7.65
Partial	0.7	98.32	1.44
	0.4	98.88	0.87
Combined	0.12	94.89	4.87
	0.42	91.67	8.09



(a) Information loss impact in CommNet.



(b) Noise impact across all models.

Figure 3: The impact of message loss on CommNet across all tested intensities (a), (b) compares the impact of three types (Gaussian noise, uniform noise and message loss) of communication disturbances at similar intensities on all models. 'Normal' (grey) signifies model performance without any communication disturbance.

level of performance degradation caused by the different types of unreliable communication and their comparative impact on model performance. This analysis will assess the model's ability to generalise under adverse communication conditions, providing a basis for comparing success rate degradation among all models tested.

Testing CommNet with varying intensities of Gaussian noise revealed an increasing, yet modest, performance degradation as noise intensity rises, see Table 3. At a noise intensity of 0.3, the success rate slightly dips, indicating a minor degradation of **0.04%**. Increasing noise to 0.5 leads to a more noticeable **0.1%** degradation and at 0.8 the success rate further decreases reflecting a **0.33%** degradation, highlighting a non-linear relationship with increasing noise intensity. This is an expected behaviour as higher noise intensities can have a compounding effect leading to more significant performance degradation. These results further highlight the resilience of CommNet to Gaussian noise, maintaining high average success rates even at increased noise levels. Uniform noise presents a different impact pattern. The degradation caused by uniform noise is consistently less than that of Gaussian noise at equivalent intensities. This suggests that CommNet is more robust to uniform noise even at higher intensities, possibly due to its uniformity which allows the model to better adjust, whereas Gaussian noise introduces more uncertainty.

Message loss significantly impacts performance, shown in Figure 3a. Partial message drops result in a degradation substantially higher compared to uniform noise at close intensity, with smaller increases in intensity leading to a significant rise in performance degradation. Whole message loss is even more detrimental. These results indicate that CommNet is significantly more sensitive to message drops than

noise, with whole message loss being particularly detrimental, suggesting that message integrity is more crucial for model success given that even small intensity can lead to a considerable performance decline. Co-occurring partial and whole message loss at a combined probability of 0.12 resulted in a performance degradation comparable to a 0.3 probability of whole message loss, suggesting that even a relatively low combined probability can have a substantial impact. A larger combined drop (at 0.42 probability) causes a degradation higher than any individual noise or message drop scenario. This indicates that the compound effect of partial and whole message loss presents a significant challenge to model performance. Overall, CommNet remains robust, maintaining high success rates even with high-intensity Gaussian and uniform noise introduced to the communication signal.

Table 4

Average success rates and percentage of degradation (in parentheses) for all models.

Condition	Average Success Rate and Degradation (%)			
	CommNet	IC3Net	GA-Comm	TARMAC
No noise	99.76	99.63	99.68	91.42
Gaussian 0.5	99.66 (0.10)	99.51 (0.12)	99.58 (0.096)	91.08 (0.37)
Uniform 0.5	99.67 (0.091)	99.44 (0.18)	99.67 (0.0092)	91.39 (0.03)
Combined 0.12	94.89 (4.87)	98.86 (0.76)	90.71 (8.96)	91.41 (0.013)

Under reliable communication conditions, CommNet and IC3Net perform comparably. With the introduction of noise, IC3Net’s performance shows a slightly steeper decline with the most notable difference in the case of combined message loss which led to a less pronounced decrease compared to CommNet, see Table 4. IC3Net’s hard attention, allowing agents to decide whether to communicate, possibly makes the model more resilient to ‘empty’ message tensors (masked by zeros in both hard attention and message drops). Similar to CommNet, IC3Net is generally robust to noise but highly vulnerable to information loss.

GA-Comm also demonstrates resilience to noise, showing only slight performance losses. However, with combined message loss, GA-Comm’s performance significantly drops, see Figure 3b. This drastic decrease contrasts with CommNet and IC3Net which maintain higher success rates under the same conditions. TARMAC reports a lower success rate under ideal communication conditions, however, degradation from noise and message drops is less pronounced. Even at combined message loss scenarios the degradation is almost negligible. TARMAC’s communication mechanism architecture possibly makes it robust against all the tested noise types, maintaining consistent performance where other models exhibit more significant losses.

4.3. Noise Comparative Impact

Comparing the impact of noise across models (see Figure 3b) reveals distinct patterns of robustness and vulnerability. Gaussian noise, with its inherent unpredictability, shows a gradual and slightly more pronounced impact on performance than uniform noise as intensity increases. Models possibly adjust better to the latter due to its uniformity which leads to slightly lower degradation (with the exception of IC3Net at 0.5 intensity) compared to Gaussian noise at the same intensity. However, both noise types show mild degradation effects on all models (Table 4). Message loss presents a different scenario. Partial message loss at a medium intensity of 0.4 causes a smaller degradation compared to whole message loss at a slightly lower intensity of 0.3 which led to significant performance drops. Increasing the combined message loss intensity from 0.12 to 0.42 (a 250% increase) resulted in a disproportional increase in degradation. For instance, in CommNet, degradation rose from 4.87% to 8.095%, representing approximately a 66% increase in degradation. This corresponds to a rate of increase in degradation of about 26.4% relative to the rate of increase in intensity. These observations suggest that noise impact depends more on the specific noise characteristics than on a simple linear progression. TARMAC’s resilience -0.013% degradation- further indicates that architecture also plays a crucial role in model response to different noise types and intensities.

5. Conclusion

This work explored the impact of different types of unreliable communication, namely noise introduced to the communication tensor and message loss, on the performance of emergent communication models in a non-signalised intersection environment. The results indicate that models are able to generalise to both Gaussian and uniform noise at low to medium intensities, suggesting that 'simple' noise can be filtered out by the inherent mechanisms of reinforcement learning. Message loss had a disproportionately greater impact on success rate compared to noise altering the conveyed information. The observed relationship between disturbance intensity and performance degradation does not appear to be linear, especially considering the disproportionate increase in degradation with whole and combined message drops. Model performance is varyingly affected by different disturbances and intensities, indicating complex dynamics rather than a straightforward linear relationship. These dynamics are likely further influenced by the specific model architecture and task difficulty.

Future work will explore the influence of training duration and extend to investigating the effects of message jumbling and delays, covering medium and hard levels of the intersection environment. This will help us firstly identify the extent to which additional training stabilises performance, especially in the presence of noise, and secondly understand the compound impact of increased task complexity and communication disruptions. Given the detrimental effects of reduced message integrity on task success, delays and jumbling could present similar challenges and potentially cause even greater performance degradation.

6. Acknowledgements

This work was supported by the Science Foundation Ireland Centre for Research Training in Advanced Networks for Sustainable Societies (ADVANCE CRT), under the Grant number 18/CRT/6222, and by the Science Foundation Ireland CONNECT Research centre Phase 2, Grant 13/RC/2077_P2. For the purpose of Open Access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

References

- [1] Y. Lin, P. Wang, M. Ma, Intelligent transportation system(its): Concept, challenge and opportunity, in: 2017 IEEE 3rd International Conference on Big Data Security on Cloud (BigDataSecurity), IEEE International Conference on High Performance and Smart Computing (Hpsc), and IEEE International Conference on Intelligent Data and Security (IDS), 2017, pp. 167–172. doi:10.1109/BigDataSecurity.2017.50.
- [2] A. A. Malikopoulos, C. G. Cassandras, Y. J. Zhang, A decentralized energy-optimal control framework for connected automated vehicles at signal-free intersections, *Automatica* 93 (2018) 244–256.
- [3] R. Hult, G. R. Campos, E. Steinmetz, L. Hammarstrand, P. Falcone, H. Wymeersch, Coordination of cooperative autonomous vehicles: Toward safer and more efficient road transportation, *IEEE Signal Processing Magazine* 33 (2016) 74–84.
- [4] M. A. Guney, I. A. Raptis, Scheduling-based optimization for motion coordination of autonomous vehicles at multilane intersections, *Journal of Robotics* 2020 (2020) 1–22.
- [5] S. Chen, J. Dong, P. Y. J. Ha, Y. Li, S. Labi, Graph neural network and reinforcement learning for multi-agent cooperative control of connected autonomous vehicles, *Computer-Aided Civil and Infrastructure Engineering* 36 (2021) 838–857. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/mice.12702>. doi:<https://doi.org/10.1111/mice.12702>. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/mice.12702>.
- [6] H. Wymeersch, G. R. de Campos, P. Falcone, L. Svensson, E. G. Ström, Challenges for cooperative its: Improving road safety through the integration of wireless communications, control, and

- positioning, in: 2015 International Conference on Computing, Networking and Communications (ICNC), 2015, pp. 573–578. doi:10.1109/ICCNC.2015.7069408.
- [7] R. Hult, G. R. Campos, P. Falcone, H. Wymeersch, An approximate solution to the optimal coordination problem for autonomous vehicles at intersections, in: 2015 American Control Conference (ACC), 2015, pp. 763–768. doi:10.1109/ACC.2015.7170826.
- [8] L. Chen, C. Englund, Cooperative intersection management: A survey, *IEEE Transactions on Intelligent Transportation Systems* 17 (2016) 570–586. doi:10.1109/TITS.2015.2471812.
- [9] A Finite State Machine Based Automated Driving Controller and its Stochastic Optimization, volume Volume 2: Mechatronics; Estimation and Identification; Uncertain Systems and Robustness; Path Planning and Motion Control; Tracking Control Systems; Multi-Agent and Networked Systems; Manufacturing; Intelligent Transportation and Vehicles; Sensors and Actuators; Diagnostics and Detection; Unmanned, Ground and Surface Robotics; Motion and Vibration Control Applications of *Dynamic Systems and Control Conference*, 2017. URL: <https://doi.org/10.1115/DSCC2017-5209>. doi:10.1115/DSCC2017-5209. arXiv:<https://asmedigitalcollection.asme.org/DSCC/proceedings-pdf/DSCC2017/58288/V002T07A002/2376129/v002t07a002-dscc2017-5209.pdf>, v002T07A002.
- [10] N. Li, H. Chen, I. Kolmanovsky, A. Girard, An explicit decision tree approach for automated driving, in: *Dynamic systems and control conference*, volume 58271, American Society of Mechanical Engineers, 2017, p. V001T45A003.
- [11] X. Lin, J. Zhang, J. Shang, Y. Wang, H. Yu, X. Zhang, Decision making through occluded intersections for autonomous driving, in: 2019 IEEE Intelligent Transportation Systems Conference (ITSC), 2019, pp. 2449–2455. doi:10.1109/ITSC.2019.8917348.
- [12] S. Gronauer, K. Diepold, Multi-agent deep reinforcement learning: a survey, *Artificial Intelligence Review* 55 (2022) 895–943.
- [13] L. Buşoniu, R. Babuška, B. D. Schutter, Multi-agent reinforcement learning: An overview, *Innovations in multi-agent systems and applications-1* (2010) 183–221.
- [14] J. Paulos, S. W. Chen, D. Shishika, V. Kumar, Decentralization of multiagent policies by learning what to communicate, in: 2019 International Conference on Robotics and Automation (ICRA), IEEE, 2019, pp. 7990–7996.
- [15] C. Yu, X. Wang, J. Hao, Z. Feng, Reinforcement learning for cooperative overtaking, in: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2019, p. 341–349.
- [16] R. Lowe, Y. WU, A. Tamar, J. Harb, O. Pieter Abbeel, I. Mordatch, Multi-agent actor-critic for mixed cooperative-competitive environments, in: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, volume 30, Curran Associates, Inc., 2017. URL: <https://proceedings.neurips.cc/paper/2017/file/68a9750337a418a86fe06c1991a1d64c-Paper.pdf>.
- [17] J. Foerster, I. A. Assael, N. De Freitas, S. Whiteson, Learning to communicate with deep multi-agent reinforcement learning, *Advances in neural information processing systems* 29 (2016).
- [18] S. Sukhbaatar, R. Fergus, et al., Learning multiagent communication with backpropagation, *Advances in neural information processing systems* 29 (2016).
- [19] W. Kim, M. Cho, Y. Sung, Message-dropout: An efficient training method for multi-agent deep reinforcement learning, in: *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 2019, pp. 6079–6086.
- [20] N. Jaques, A. Lazaridou, E. Hughes, C. Gulcehre, P. Ortega, D. Strouse, J. Z. Leibo, N. De Freitas, Social influence as intrinsic motivation for multi-agent deep reinforcement learning, in: K. Chaudhuri, R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, PMLR, 2019, pp. 3040–3049. URL: <https://proceedings.mlr.press/v97/jaques19a.html>.
- [21] D. Simões, N. Lau, L. P. Reis, Multi-agent actor centralized-critic with communication, *Neurocom-*

puting 390 (2020) 40–56.

- [22] A. Das, T. Gervet, J. Romoff, D. Batra, D. Parikh, M. Rabbat, J. Pineau, TarMAC: Targeted multi-agent communication, in: K. Chaudhuri, R. Salakhutdinov (Eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, PMLR, 2019, pp. 1538–1546. URL: <https://proceedings.mlr.press/v97/das19a.html>.
- [23] Y. Niu, R. R. Paleja, M. C. Gombolay, Multi-agent graph-attention communication and teaming., in: *AAMAS*, 2021, pp. 964–973.
- [24] A. Singh, T. Jain, S. Sukhbaatar, Learning when to communicate at scale in multiagent cooperative and competitive tasks, in: *International Conference on Learning Representations*, 2018.
- [25] R. Wang, X. He, R. Yu, W. Qiu, B. An, Z. Rabinovich, Learning efficient multi-agent communication: An information bottleneck approach, in: *International Conference on Machine Learning*, PMLR, 2020, pp. 9908–9918.
- [26] T.-Y. Tung, S. Kobus, J. P. Roig, D. Gündüz, Effective communications: A joint learning and communication framework for multi-agent reinforcement learning over noisy channels, *IEEE Journal on Selected Areas in Communications* 39 (2021) 2590–2603.
- [27] J. S. P. Roig, D. Gündüz, Remote reinforcement learning over a noisy channel, in: *GLOBECOM 2020-2020 IEEE Global Communications Conference*, IEEE, 2020, pp. 1–6.
- [28] Y. Liu, W. Wang, Y. Hu, J. Hao, X. Chen, Y. Gao, Multi-agent game abstraction via graph attention neural network, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2020, pp. 7211–7218.
- [29] T. Eccles, Y. Bachrach, G. Lever, A. Lazaridou, T. Graepel, Biases for emergent communication in multi-agent reinforcement learning, *Advances in neural information processing systems* 32 (2019).
- [30] M. Tan, Multi-agent reinforcement learning: Independent vs. cooperative agents, in: *Proceedings of the tenth international conference on machine learning*, 1993, pp. 330–337.
- [31] F. Qureshi, D. Terzopoulos, Smart camera networks in virtual reality, *Proceedings of the IEEE* 96 (2008) 1640–1656.
- [32] K. Zhang, Z. Yang, H. Liu, T. Zhang, T. Basar, Fully decentralized multi-agent reinforcement learning with networked agents, in: J. Dy, A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, PMLR, 2018, pp. 5872–5881. URL: <https://proceedings.mlr.press/v80/zhang18n.html>.
- [33] J. Jiang, Z. Lu, Learning attentional communication for multi-agent cooperation, *Advances in neural information processing systems* 31 (2018).
- [34] H. Mao, Z. Zhang, Z. Xiao, Z. Gong, Y. Ni, Learning agent communication under limited bandwidth by message pruning, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2020, pp. 5142–5149.
- [35] C. Luo, X. Liu, X. Chen, J. Luo, Multi-agent fault-tolerant reinforcement learning with noisy environments, in: *2020 IEEE 26th International Conference on Parallel and Distributed Systems (ICPADS)*, 2020, pp. 164–171. doi:10.1109/ICPADS51040.2020.00031.
- [36] C. E. Shannon, A mathematical theory of communication, *The Bell system technical journal* 27 (1948) 379–423.
- [37] A. Mostaani, O. Simeone, S. Chatzinotas, B. Ottersten, Learning-based physical layer communications for multiagent collaboration, in: *2019 IEEE 30th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, IEEE, 2019, pp. 1–6.
- [38] J. S. P. Roig, D. Gündüz, Remote reinforcement learning over a noisy channel, in: *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, 2020, pp. 1–6. doi:10.1109/GLOBECOM42002.2020.9322408.
- [39] J. Blumenkamp, A. Prorok, The emergence of adversarial communication in multi-agent reinforcement learning, in: J. Kober, F. Ramos, C. Tomlin (Eds.), *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, PMLR, 2021, pp. 1394–1414. URL: <https://proceedings.mlr.press/v155/blumenkamp21a.html>.
- [40] A. Mostaani, O. Simeone, S. Chatzinotas, B. Ottersten, Learning-based physical layer communications for multiagent collaboration, in: *2019 IEEE 30th Annual International Symposium on*

Personal, Indoor and Mobile Radio Communications (PIMRC), 2019, pp. 1–6. doi:10.1109/PIMRC.2019.8904190.

- [41] Y. Niu, R. Paleja, M. Gombolay, Multi-agent graph-attention communication and teaming, in: Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '21, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2021, p. 964–973.