

Effective Communication with Dynamic Feature Compression

Pietro Talli, *Student Member, IEEE*, Francesco Pase, *Graduate Student Member, IEEE*,
Federico Chiariotti, *Member, IEEE*, Andrea Zanella, *Senior Member, IEEE*, and Michele Zorzi, *Fellow, IEEE*

Abstract—The remote wireless control of industrial systems is one of the major use cases for 5G and beyond systems: in these cases, the massive amounts of sensory information that need to be shared over the wireless medium may overload even high-capacity connections. Consequently, solving the *effective communication* problem by optimizing the transmission strategy to discard irrelevant information can provide a significant advantage, but is often a very complex task. In this work, we consider a prototypal system in which an observer must communicate its sensory data to a robot controlling a task (e.g., a mobile robot in a factory). We then model it as a remote Partially Observable Markov Decision Process (POMDP), considering the effect of adopting semantic and effective communication-oriented solutions on the overall system performance. We split the communication problem by considering an ensemble Vector Quantized Variational Autoencoder (VQ-VAE) encoding, and train a Deep Reinforcement Learning (DRL) agent to dynamically adapt the quantization level, considering both the current state of the environment and the memory of past messages. We tested the proposed approach on the well-known CartPole reference control problem, obtaining a significant performance increase over traditional approaches.

Index Terms—Effective communication, Networked control, Semantic communication, Information bottleneck

I. INTRODUCTION

THE main goal of classical communication theory is to build reliable systems for the accurate transmission of arbitrary data through a constrained communication channel while using as few symbols as possible. However, in the preface to Shannon's seminal work [1], Warren Weaver already envisioned two more complex levels of communication beyond the simple transmission of bits. Classical communications are then included in Level A, or the *technical problem*, which concerns itself with the accurate and efficient transmission of arbitrary raw data. Level B, or the *semantic problem*, is to find the best way to convey the meaning of the message, even when irrelevant details are lost or misunderstood, while Level C, also called the *effectiveness problem*, deals with the resulting behavior of the receiver [2]: as long as the receiver takes the optimal decision, the effectiveness problem is solved, regardless of the quality of the received information.

Pietro Talli (corresponding author, pietro.talli@phd.unipd.it), Francesco Pase (pasefrance@dei.unipd.it), Federico Chiariotti (chiariot@dei.unipd.it), Andrea Zanella (zanella@dei.unipd.it) and Michele Zorzi (zorzi@dei.unipd.it) are with the Department of Information Engineering, University of Padova, 35131 Padua, Italy.

This work was supported by the European Union under the Italian National Recovery and Resilience Plan of NextGenerationEU, under the partnership on "Telecommunications of the Future" (PE0000001 - program "RESTART") and the "SoE Young Researchers" grant REDIAL.

The Level B and C problems are tightly intertwined, as defining the meaning of a message is often related to the intentions of the receiver. The difference between Level B and C communication lies in the way the behavioral model of the receiver is exploited by the transmitter to encode the message to be delivered. In Level B, the transmitter is not influenced by the receiver's policy, and indeed might consider a passive receiver such as a monitoring application, while in Level C, the sender explicitly takes the receiver's policy into account, crafting the message so as to drive the receiver's actions.

While the Level B and C problems attracted limited attention for decades, the explosion of Industrial Internet of Things (IIoT) systems has drawn the research and industrial communities toward semantic and effective communication [3], optimizing remote control processes under severe communication constraints beyond Shannon's limits on Level A performance [2]. In particular, the effectiveness problem is highly relevant to robotic applications, in which independent mobile robots, such as drones or rovers, must operate based on information from remote sensors. In this case the sensors and the cameras act as the transmitter in a communication problem, while the robot is the receiver: by solving the Level C problem, the sensors can transmit the information that best directs the robot's actions toward the optimal policy [4]. We can also consider a case in which the robot is the transmitter, while the receiver is a remote controller, which must get the most relevant information to decide the control policy [5].

The rise of communication metrics that take the content of the message into account, such as the Value of Information (VoI) [6], represents an attempt to approach the problem in practical scenarios, and analytical studies have exploited information theory to define a semantic accuracy metric and minimize distortion [7]. In particular, *information bottleneck* theory [8] has been widely used to characterize Level B optimization [9]. However, translating a practical system model into a semantic space is a nontrivial issue, and the semantic problem is a subject of active research [10], [11]. The effectiveness problem is even more complex, as it implicitly depends on estimating the effect of communication distortion on the control policy and, consequently, on its performance [12]. While the effect of simple scheduling policies is relatively easy to compute [13], and linear control systems can be optimized explicitly [14], realistic control tasks are highly complex, complicating an analytical approach to the Level C problem. Pure learning-based solutions that consider communication as an action in a multi-agent DRL problem, such as emergent communication, also have limitations [15],

as they can only deal with very simple scenarios due to significant convergence and training issues. In some cases, the information bottleneck approach can also be exploited to determine state importance [16], but the existing literature on optimizing Level C communication is very sparse and limited to simpler scenarios [17]. Another possible approach to the remote tracking of Markov sources is addressed in zero-delay coding theory [18]. However, this theory considers the error on the hidden state estimate as the objective of the optimization, which is similar to what we could consider a semantic (or Level B) approach. In this work, we show that it is possible to optimize the system with respect to other metrics which cannot be explicitly derived, such as cumulative rewards in DRL: by accepting a higher distortion at the semantic level when it is not relevant to the task, we can further reduce the required bitrate without sacrificing the control performance.

In this work, we consider a dual model which combines concepts from DRL and semantic source coding: we consider an ensemble of VQ-VAE models [19], each of which learns to represent observations using a different codebook. A DRL agent can then select the codebook to be used for each transmission, controlling the trade-off between accuracy and compression. Depending on the receiver's task, the reward for the agent DRL can be adjusted to solve the Level A, B, and C problems, optimizing the performance for each specific task. In order to test the performance of the proposed framework in a relevant example scenario, we consider the well-known CartPole problem, whose state can be easily converted into a semantic one, as its dynamics depend on a limited set of physical quantities. The problem we selected is purposefully simple, as this allows for a better explainability and an easier training, but the solution is not limited to the CartPole problem, and can be adapted to more complex tasks. The main contributions of this paper are then given by the following:

- We model a remote-control system as a remote POMDP problem and present an efficient solution for learning effective communication through the dynamic compression of learnable features;
- We show that dynamic codebook selection outperforms static strategies for all three levels, and that considering the Level C task can significantly improve the control performance without increasing the bitrate;
- We adopt an explainability framework to understand the choices of the agent in this simple problem, and verify that the Level C dynamic compression captures the receiver's uncertainty in the state estimation and its impact on the expected reward, transmitting only when necessary.

We remark that the dynamic codebook selection policy is not limited to the VQ-VAE ensemble we consider, but is a general technique that can be applied to any compression algorithm with adaptable quality parameters, helping to deliver more accurate information when it is relevant to do so. The results and policy analysis lead to significant insights for the design of communication strategies for remote control. A partial version of this work was presented at the IEEE INFOCOM WiSARN 2023 workshop [20]. This paper includes a more complete

theoretical characterization of the problem, as well as an updated learning architecture, additional results, and an in-depth analysis of the DRL-based dynamic compression policy.

The rest of the paper is organized as follows: first, we analyze the state of the art on semantic and effective communication in Sec. II. Sec. III then presents the general system model and the three levels of communication we consider. We then describe the dynamic feature compression solution in Sec. IV, which is evaluated by simulation in Sec. V. Finally, Sec. VI concludes the paper.

II. RELATED WORK

The specific requirements of distributed and remotely controlled systems have focused the research community's attention towards communication systems that must provide updated information to enable real-time high-level tasks such as inference, tracking, or control. Although metrics such as Age of Information (AoI) [6] represent a major improvement with respect to latency and packet loss, they are still limited, as they assume that the quality of the information available at the receiver degrades deterministically with time, most commonly (but not necessarily [21]) in a linear fashion. However, more sophisticated systems can also take into consideration the current state of the system in order to decide whether and when to update the status of the receiver. Metrics such as Urgency of Information (UoI) [14] and VoI [22] incorporate state information in their definition and are thus aware of the intrinsic value of potential updates. Other context-aware indices to measure the nonlinear time-varying importance and the non-uniform context dependence of the status information have also been proposed [14]. The authors of [23] considered a system in which a transmitter monitors the status of a system and updates the controller, providing status information. Then a constrained Markov decision process (MDP) is formulated to minimize the cost of actuation and simultaneously guarantee a target communication rate. Both works show significant improvements with respect to other metrics such as AoI.

At the same time, the development of learning-based coding schemes has allowed communication system designers to move beyond packet error as the key coding performance metric, exploiting semantic considerations. Joint source-channel coding for wireless image transmission is implemented in [24], [25], and the encoder-decoder pair is parameterized by a neural network (NN), whose architecture may vary. This approach can be used to maintain the semantic information contained in the transmitted data, while improving the compression performance.

Semantic information at the receiver can be used to solve different tasks. Effective communication [12] can be seen as an extension of this, in which the task involves the receiver taking actions and possibly altering the information that the transmitter is communicating. Effective communication differs from semantic communication mostly because the "semantic" content which has to be preserved in the communicated messages is not explicit. Moreover, control tasks have a temporal component that must be taken into account, as investigated in [12]. The scenario considered in the work is a two-agent

POMDP in which one agent communicates and the other agent interacts with the environment, using DRL to solve the joint problem and encoding the information. A distributed perception scenario, in which multiple sensors communicate to a single robot, is considered in [26] and solved using multi-agent reinforcement learning (MARL), showing that joint training improves the performance of the system, particularly when communication is severely constrained. While past works aimed at specific scenarios and objectives, this paper proposes a novel DRL approach that combines status updates with an adaptive coding scheme and can be easily adapted to operate at any of the three levels of communication (A, B, or C).

III. SYSTEM MODEL

The recent interest in semantic and effective communication from the research community has driven the development of a wide array of models and conceptualizations, as highlighted in the previous section. At the highest level of abstraction, our purpose is to define a model in which effective communication is meaningful, and the differences between the three problems in Weaver's formulation become clear.

Let us then consider a simple example: we have a remote actuator performing a control task, while a camera observes the results and transmits its observation through a wireless channel. The actuator might have its own sensors, but it relies on the video feed to improve its performance and maintain stable and efficient control. The classical Level A approach to the problem would be to compress the video as efficiently as possible, minimizing the reconstruction error on frames by using an appropriate codec. The difference between Level A and Level B solutions is then obvious: the former encodes new frames so that the reconstruction fidelity is preserved, while the latter maps elements in the frame to their importance when estimating the physical state of the system. In some control applications, the state can be defined trivially, while in others it may be more complex, but in general, the translation of the video to the state space is unaffected by irrelevant information (such as, e.g., movements in the background).

If we consider Level C, we target control performance directly, and thus further restrict the definition of relevant information: while Level B concerns itself with estimating the system state correctly, a Level C solution only considers errors in the state estimation when they cause performance drops. If the control action is the same over a wide set of states, accuracy then becomes unnecessary, as the actuator only needs a rough estimate of the state to decide what to do; the same happens if there are multiple actions with almost equivalent performance, i.e., if the optimality gap caused by imperfect information remains small.

These natural observations represent the core concepts of effective communication, but implementing them in practical systems is often complex as actions have long-term consequences, and state estimates are based on a history of observations, so that transmitting a message may affect future performance in complex ways. In the following, we provide an analytical framework using the *remote POMDP* approach to objectively evaluate these choices and implement a solution

TABLE I: Main notation and definitions.

Symbol	Definition
\mathcal{S}	Set of system states
\mathcal{A}	Set of feasible actions
\mathcal{O}	Set of observations
P	State transition probability function
ω	Observation function
R	Reward function
γ	Discount factor
h	History of stochastic observations
$h^{(r)}$	History of received messages at the robot
π	Policy
G	Expected cumulative discounted reward
$\Phi(\cdot)$	Space of probability distributions over a set
ξ	Belief distribution over the state space \mathcal{S}
$\xi^{(o)}$	Belief distribution at the observer
$\xi^{\text{pri.}(r)}$	Prior belief distribution at the robot
$\xi^{(r)}$	Belief distribution at the robot
\mathcal{M}	Set of messages
Λ	Encoding function
m	Message communicated
$\ell(m)$	Length of message m
ζ	Vector quantizer
\mathcal{P}	Picture space
β	Communication cost

for effective communication in cyber-physical systems. We will denote random variables with capital letters, their possible values with lowercase letters, and sets with calligraphic or Greek capitals. Table I reports the main symbols we introduce in the following sections for the reader's convenience.

A. POMDP Definition and Solution

In the infinite horizon POMDP formulation [27], one agent needs to optimally control a stochastic process defined by a tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{O}, P, \omega, R, \gamma \rangle$, where \mathcal{S} represents the set of system states, \mathcal{A} is the set of feasible actions, and \mathcal{O} is the observation set. The function $P : \mathcal{S} \times \mathcal{A} \rightarrow \Phi(\mathcal{S})$, where $\Phi(\cdot)$ represents the space of probability distributions over a set, gives the state transition probability function. We denote the conditional probability distribution of the next state, given the current state and the selected action, as $P(s'|s, a) = \Pr[S_{t+1} = s' | S_t = s, A_t = a]$. Then, $\omega : \mathcal{S} \rightarrow \Phi(\mathcal{O})$ denotes the conditional probability of the observation O_t given the system state $s \in \mathcal{S}$, so that $\omega(o|s) = \Pr[O_t = o | S_t = s]$, is the probability of receiving the observation $o \in \mathcal{O}$, given that the system is in state $s \in \mathcal{S}$. Finally, function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ represents the expected reward received by the agent when taking action a in state s , denoted as $R(s, a)$, and the scalar $\gamma \in [0, 1)$ is a discount factor used to compute the long-term reward. We notice that functions P , R , and ω do not depend on the time instant t , thus focusing on homogeneous Markov processes.

The POMDP proceeds in discrete steps indexed by t : at each step t , the agent can infer the system state s_t only from the partial information given by the history of stochastic observations $h_t = (o_t, o_{t-1}, \dots, o_1) \in \mathcal{O}^t$. Based on these observations, and on its policy $\pi : \mathcal{O}^t \rightarrow \Phi(\mathcal{A})$, which outputs a probability distribution over the action space for each possible observation history, the agent interacts with the system by selecting an action $A_t \sim \pi(h_t)$. The sampled action a_t is

then performed in the real environment, whose hidden state s_t is unknown to the agent, which then receives a feedback from the environment in the form of a (potentially stochastic) reward r_t , with expected value $R_t = R(s_t, a_t)$.¹ The goal for the agent is then to optimize its policy π to maximize the expected cumulative discounted reward $G = \mathbb{E} \left[\sum_t \gamma^t R_t \right]$. Having an optimal policy is equivalent to knowing the optimal state-action values, also known as Q -values, and taking action

$$a_t = \pi(h_t) = \arg \max_{a \in \mathcal{A}} Q(h_t, a),$$

where $Q(h_t, a) = \mathbb{E} \left[\sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_{\tau} | h_t, a \right]$ is the expected cumulative reward starting from (h_t, a) .

However, considering the full history of observations makes the solution highly complex, as the length of h_t is potentially unbounded. We then define an estimator $\xi : \mathcal{O}^t \rightarrow \Phi(S)$, which outputs the *a posteriori* belief distribution over the state space. We can then recast the original POMDP as a standard MDP, whose state space is $S' = \Phi(S)$, i.e., the space of possible belief distributions. Solving the POMDP in this modified belief space has been proved to be optimal in [28].

The policy over this modified MDP is then $\pi : S' \rightarrow \mathcal{A}$, which can be optimized using standard tools [29]. We can also compute the new transition probability and expected reward as in [28]. Given $\xi_t(s) = \Pr(S_t = s | h_t)$, which represents the maximum likelihood estimate of the state given all the history of observations h_t , and $A_t = a$, we define the *a priori* belief over the state at time $t + 1$ as

$$\begin{aligned} \xi_{t+1}^{\text{pri}}(s | \xi_t, a) &= \Pr[S_{t+1} = s | \xi_t, A_t = a] \\ &= \sum_{s' \in S} \xi_t(s') P(s | s', a). \end{aligned}$$

The *a posteriori* belief ξ_{t+1} , which also includes the new observation o_{t+1} , can then be obtained by performing a Bayesian update using the *a priori* belief as a prior:

$$\begin{aligned} \xi_{t+1}(s | \xi_t, a, o) &= \Pr[S_{t+1} = s | \xi_t, A_t = a, O_{t+1} = o] \\ &= \Pr \left[S_{t+1} = s | \xi_{t+1}^{\text{pri}}(s | \xi_t, a), O_{t+1} = o \right] \\ &= \frac{\omega(o | s) \xi_{t+1}^{\text{pri}}(s | \xi_t, a)}{\sum_{s' \in S} \xi_{t+1}^{\text{pri}}(s' | \xi_t, a) \omega(o | s')}. \end{aligned} \quad (1)$$

These update equations allow us to compute the modified transition probability matrix P' for the belief MDP, which is then defined by the tuple $\langle \Phi(S), \mathcal{A}, P', R, \gamma \rangle$.

B. The Remote POMDP

In this paper, we consider a variant of the POMDP, that we define *remote POMDP*, in which two agents are involved in the process. The first agent, i.e., the *observer*, receives observation $O_t \in \mathcal{O}$, and needs to convey such information to a second agent, i.e., the *robot*, through a constrained communication channel, which limits the number of bits the observer can send reliably. Consequently, the amount of information the observer can send to the robot is limited. The robot then chooses

and takes an action in the physical environment. This system can formalize many control problems in future IIoT systems, as sensors and actuators may potentially be geographically distributed, and the amount of information they can exchange to accomplish a task is limited by the shared wireless medium, which has to be allocated to the many devices installed in the factory, as well as by the energy limitations of the sensors. Similar systems have been analyzed in [12], [23]. We will now analyze the problems for the two agents, considering a case in which communication and control are designed separately. Joint control and communication approaches [26] can outperform separate approaches by tuning the two agents' policies to each other, but they introduce additional training complexity, and will not be considered in this work. In the following, we will refer to variable x related to the robot as $x^{(r)}$, while the corresponding variable on the observer side will be denoted by $x^{(o)}$.

1) *The Robot-Side POMDP*: We denote the message communicated to the robot at time step t as $m_t \in \mathcal{M}$. The set of possible messages \mathcal{M} forms the set of observations that are available to the robot, and the history of these observations is given by $h_t^{(r)} = \{m_t, \dots, m_1\}$, which is the sequence of messages received up to time t . We can then see the robot as an agent with its own POMDP, in which the observations are filtered by both the partial knowledge of the observer and the further distortion produced by the fact that these observations are encoded and communicated through a constrained channel. The robot-side POMDP is then defined by the tuple $\langle S, \mathcal{A}, \mathcal{M}, P, \pi^{(o)}, R, \gamma \rangle$, as observations depend on the observer's policy.

The message transmitted from the observer to the robot modifies the belief distribution over the next state as a Bayesian update. Let us define the distribution over the current state, given that the message m_t has been received, as $\xi_t^{(r)}$. For example, if the communicated message m_t contains the correct state $S_t = s$, the belief distribution becomes deterministic, i.e., $\xi(s' | m_t) = \delta_{s, s'}$, where $\delta_{m, n}$ is the Kronecker delta function, equal to 1 if the two arguments are the same and 0 otherwise. Ideally, an intelligent observer will allocate more communication resources and thus provide more precise messages if the *a priori* distribution of the robot is far from the one estimated by the observer. The modified MDP is then defined by the tuple $\langle \Phi(S), \mathcal{A}, P^{(r)}, R, \gamma \rangle$, where $P^{(r)}$ represents the Bayesian update function. According to the previous notation, we can express the optimal action at time $t + 1$ as

$$a_{t+1} = \arg \max_{a \in \mathcal{A}} Q \left(\xi_t^{(r)}, a \right),$$

where $\xi_t^{(r)}$ is the current belief at the robot side (after message m_t is received). The robot's reward is simply given as the reward of the original POMDP, i.e., the control performance in the environment. The optimal policy can be reached by using standard DRL tools.

2) *The Observer-Side POMDP*: On the other side, the observer needs to encode its belief $\xi_t^{(o)}$ in a message $m_t \in \mathcal{M}$ and transmit it. We can then consider the observer-side POMDP, in which the action set corresponds to the set of

¹As the reward signal is only provided during training, it cannot be used to infer the value of s_t .

messages \mathcal{M} and the state space is represented by the belief from the observed results. The tuple defining this POMDP is $\langle \mathcal{S}, \mathcal{M}, \mathcal{O}, P, \omega, R^{(o)}, \gamma \rangle$. We assume that the observer knows the robot's policy, i.e., it can know the actions that the robot takes in the environment and use them to improve its estimate of the state. This can also be accomplished if the robot transmits the actions it takes as feedback to the observer. As described above for the robot-side problem, we can transform this POMDP into the belief MDP given by $\langle \Phi(\mathcal{S}) \times \Phi(\mathcal{S}), \mathcal{M}, P^{(o)}, R^{(o)}, \gamma \rangle$, where $P^{(o)}$ is the Bayesian update given in (1). We highlight that the observer needs to keep track of both its own and the robot's belief, as the effectiveness of communication depends on the difference between the two, and the state of the observer is given by $\langle \xi_t^{(o)}, \xi_t^{\text{pri},(r)} \rangle$.

The objective of the observer is to minimize channel usage, i.e., communicate as few bits as possible, while maintaining the highest possible performance in the control task: the expected reward $R^{(o)}$ depends on both components. We then consider a simple linear combination approach: if it transmits message m , whose length in bits is $\ell(m)$, the observer then gets a penalty $\beta\ell(m)$, where $\beta \in \mathbb{R}^+$ is a cost parameter. In order to optimize its policy, the observer also needs to have a way to gauge the *value* of information, which is a complex problem: information theory, and in particular rate-distortion theory, have provided the fundamental limits when optimizing for the technical problem, i.e., Level A, where the goal is to reconstruct the source signals with the highest fidelity with a given maximum rate [30], or, conversely, encode it with the minimum rate to achieve a given reconstruction fidelity. We will discuss the definition of VoI in the following sections. More complex modeling choices for the transmission cost are also possible, e.g., considering energy constraints for a sensor with energy harvesting capabilities, but are beyond the scope of this work.

As the complexity of the problem is massive, we restrict ourselves to a smaller action space by making a simplifying assumption, which allows us to separate the problem: the observer does not transmit the entire belief distribution, which may be implicit, but rather the observation O_t . We then consider the encoding function $\Lambda : \mathcal{O} \times \Phi(\mathcal{S}) \rightarrow \mathcal{M}$, which will generate a message $M_t = \Lambda(O_t | \xi_t^{(o)})$ to be sent to the robot at each step t .

C. Observer Reward in Remote POMDPs

The first and simplest way to solve the remote POMDP problem is to blindly apply a standard Level A rate-distortion metric to compress the sensor observations into messages to be sent to the agent. As an example, in the CartPole problem analyzed in this work (see Sec. V), one sensor observation is given by two consecutive 2D camera acquisitions. The observer's policy is then independent from the robot's task, and can be computed separately. The Level A reward function $R_A^{(o)}$ is then given as follows:

$$R_A^{(o)} \left(\langle \xi_t^{(o)}, \xi_t^{\text{pri},(r)} \rangle, m \right) = -d_A(o_t, \hat{o}_t) - \beta\ell(m). \quad (2)$$

In the CartPole case, a natural distortion metric is the image Peak Signal to Noise Ratio (PSNR), an image quality metric proportional to the logarithm of the normalized Mean Square Error (MSE) between the images. Naturally, encoding the observation with a higher precision will require more bits, as the set of messages needs to be bigger.

The Level B problem considers the projection of the raw observations into a significantly smaller *semantic space*, over which we measure distortion using function d_B , explicitly capturing the error over the needed physical system information, e.g., the angular position and velocity of the pole in the CartPole problem. The Level B reward function $R_B^{(o)}$ is then given as follows:

$$R_B^{(o)} \left(\langle \xi_t^{(o)}, \xi_t^{\text{pri},(r)} \rangle, m \right) = -d_B \left(\xi_t^{(o)}, \xi_t^{(r)} \right) - \beta\ell(m). \quad (3)$$

In our CartPole case, this may be simply represented by the MSE between the best estimate of the state at the transmitter and the receiver.

Finally, we can consider the Level C system. In this case, the distortion metric is not needed, as the control performance can be used directly, and the reward $R_C^{(o)}$ is:

$$R_C^{(o)} \left(\langle \xi_t^{(o)}, \xi_t^{\text{pri},(r)} \rangle, m \right) = R \left(\xi_t^{(r)}, \pi^{(r)} \left(\xi_t^{(r)} \right) \right) - \beta\ell(m). \quad (4)$$

The VoI of message m , $V \left(\xi_t^{\text{pri},(r)}, m \right)$, can then be given by the difference between the expected performance of the robot with and without this information:

$$V \left(\xi_t^{\text{pri},(r)}, m \right) = Q \left(\xi_t^{(r)}, \pi^{(r)} \left(\xi_t^{(r)} \right) \right) - Q \left(\xi_t^{\text{pri},(r)}, \pi^{(r)} \left(\xi_t^{\text{pri},(r)} \right) \right). \quad (5)$$

Thus, the optimal Level C observer policy $\pi_C^{(o)}$ will balance the trade-off between the performance at the receiver and the communication cost not only in the current time step, but also in the long term. This foresighted behavior is essential when considering that the belief distributions incorporate the memory of previously received messages. Providing information that does not improve the expected reward in the next step might still be worth the cost if it allows the robot to improve its estimate, reducing the need for future communication.

IV. PROPOSED SOLUTION

In this section, we introduce the architecture we used to represent Λ , the VQ-VAE, and discuss the remote POMDP solution. As the VQ-VAE model is not adaptive, we consider an ensemble model with different quantization levels, limiting the choice to *which* VQ-VAE model to use for the transmission. In our scheme, the observer is composed by the sensor (i.e., the camera that captures the system images), the ensemble of VQ-VAE encoders, each of which compresses the images at a different level, and the DRL agent that learns which encoder to select based on the current sensor's observation and the purpose of the considered performance metric at the receiver. As we mentioned, directly learning the encoding is highly complex, with a vast action space, and techniques such as emergent communication that learn it explicitly are limited to scenarios with very simple tasks and immediate rewards. By

restricting the problem to a smaller action space, we find a potentially sub-optimal solution, but we can deal with much more complex problems.

A. Deep VQ-VAE Encoding

In order to represent the encoding function Λ , and to restrict the observer-side POMDP to a more manageable action space, the observer exploits the VQ-VAE architecture introduced in [19]. The VQ-VAE is built on top of the more common Variational Autoencoder (VAE) model, with the additional feature of finding an optimal discrete representation of the latent space. The VAE is used to reduce the dimensionality of an input vector $X \in \mathbb{R}^I$, by mapping it to a stochastic latent representation $Z \in \mathbb{R}^K \sim q_\kappa(Z|X)$, where $K < I$. The stochastic encoding function $q_\kappa(Z|X)$ is a parameterized probability distribution represented by a neural network with parameter vector κ . To find optimal latent representations Z , the VAE jointly optimizes a decoding function $p_\theta(\hat{X}|Z)$ that aims to reconstruct X from a sample $\hat{X} \sim p_\theta(\hat{X}|Z)$. In this way, the parameter vectors κ and θ are usually jointly optimized to minimize the distortion $d(X, \hat{X})$ between the input and its reconstruction, given the constraint on Z , while reducing the distance between $q_\kappa(Z|X)$, and some prior $q(Z)$ [31] used to impose some structure or complexity budget.

However, in practical scenarios, one needs to digitally encode the input X into a *discrete* latent representation for digital transmission, mapping it to a dictionary \mathcal{M} . To do this, the VQ-VAE quantizes the latent space by using N codewords $z_1, \dots, z_N \in \mathbb{R}^K$ in the continuous latent space, forming a dictionary with N entries. Moreover, to better represent complex inputs, the VQ-VAE quantizes the latent representation Z as the concatenation of F features $f(X)$ of the input, each of which uses a set of N possible codewords. We denote the set containing all the N^F possible concatenated discrete features with $\mathcal{M}(F, N)$, as it represents the set of all possible messages that the observer can use to convey to the robot the information on the observation O . The peculiarity of the VQ-VAE architecture is that it jointly optimizes the codewords in $\mathcal{M}(F, N)$ together with the stochastic encoding and decoding functions q_κ and p_θ , instead of simply applying fixed vector quantization on top of the learned continuous latent variable Z . When the communication budget is fixed, i.e., the observer can transmit L bits in each round, the protocol to solve the remote POMDP is rather simple: first, the observer trains the VQ-VAE with $N(F, L) = 2^{F-1L}$, mapping each codeword to a string of $\frac{L}{F}$ bits, to minimize the technical, semantic, or effective distortion d_α , depending on the problem; then, at each step t , the observer computes $\hat{m} \sim q_\kappa(\cdot|o_t)$, and finds $m_t = \arg \min_{m \in \mathcal{M}(F, N(F, L))} \|m - \hat{m}\|_2$. The message m_t is sent to the robot, which can optimize its decision accordingly.

B. Dynamic Feature Compression

We can then consider the architecture shown in Fig. 1, consisting of a set of VQ-VAEs $\mathcal{V} = \{\zeta_\emptyset, \zeta_1, \dots, \zeta_V\}$, where each VQ-VAE ζ_v compresses each feature using v bits, i.e., with $N = 2^v$ quantized codewords. The length of a message

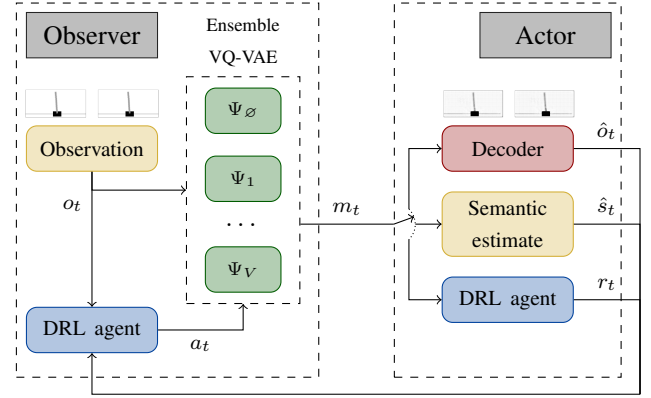


Fig. 1: Dynamic feature compression architecture.

encoded with ζ_v is then vF bits. We also include a null action ζ_\emptyset , which corresponds to not transmitting anything. As we only consider the communication side of the problem, the actor is trained beforehand using the messages with the finest-grained quantization, which are compressed with the VQ-VAE ζ_V with the largest codebook. This choice ensures that the actor can deal with finer-grained inputs, while still being robust to lower-precision features. The robot can then perform three different tasks, corresponding to the three communication problems: it can decode the observation (Level A) with the highest possible accuracy, using the decoder part of the VQ-VAE architecture; it can estimate the hidden state (Level B) using a supervised learning solution; or it can perform a control action based on the received information and observe its effects (Level C).

In all three cases, the dynamic compression is performed by the observer, based on the feedback from the robot. The observer side of the remote POMDP, whose reward is given in (2)-(4), is restricted to the choice of ζ_v , i.e., to selecting one of the possible codebooks learned by each VQ-VAE in the ensemble model, or to avoid any transmission. As we described in the previous section, the type of reward depends on the communication problem that the observer is trying to solve: at Levels A and B, the observer aims at minimizing distortion in the observation and semantic space, respectively. At Level C, the objective is to maximize the robot's reward. We remark that the policy of the receiver affects the transmitter decisions only at Level C. In fact, the goal of Level B is to extract the semantic information related to the physical process that models the dynamics of the system from the images. The optimization of the Observer (transmitter) is then aimed at minimizing the reconstruction error on this semantic estimate, but it does not consider the decision of the Actor (receiver). Conversely, Level C optimization does consider the influence of the observed physical state of the system on the Actor, reducing the redundant or irrelevant information with respect to the agent's decision.

In all three cases, *memory* is important: representing snapshots of the physical system in consecutive instants, subsequent observations have high correlations, and the robot can glean a significant amount of information from past messages. This is an important advantage of dynamic compression, as it can adapt messages to the estimated knowledge of the receiver.

While the observer adapts its transmissions to the robot’s task, the robot’s algorithms are fixed. They could themselves be adapted to the dynamic compression strategy, but this joint training is significantly more complex, and we consider it as a possible extension of this work.

C. RL implementation

There are two policies in the considered system, one for the observer and one for the robot and in both cases the policies are learned through the Actor Critic algorithm. This means that an agent learns a parametric policy π_λ and a Q -value estimator. Both the policy and the Q -values are neural networks. In order to take into account the past observations the two networks share a Long Short-Term Memory (LSTM) layer which estimates a latent state, which is then given as input to both the policy and the values estimator. This architecture avoids explicitly modeling the belief distribution which may be complicated to treat in continuous settings like the one considered in this work. This practical choice is also useful to avoid decoding the latent features discovered by the VQ-VAE back in the observation space \mathcal{O} or in the physical state space \mathcal{S} , increasing the potential for errors. Indeed, the quantized features communicated with the message m_t contain a structured representation of the observation space which can be used effectively by an LSTM to estimate the *true state*. The training algorithm is the standard Advantage Actor Critic (A2C), but the replay buffer is appropriately modified to take into account the history of previously received messages.

V. SIMULATION SETTINGS AND RESULTS

The underlying use case analyzed in this work is the well-known CartPole problem, as implemented in the *OpenAI Gym* library.² In this problem, a pole is installed on a cart, and the task is to control the cart position and velocity to keep the pole in equilibrium. The physical state of the system is fully described by the cart position x_t and velocity \dot{x}_t , and the pole angle ψ_t and angular velocity $\dot{\psi}_t$. Consequently, the true state of the system is $s_t = (x_t, \dot{x}_t, \psi_t, \dot{\psi}_t)$, and the semantic state space is $\mathcal{S} \subset \mathbb{R}^4$ (because of physical constraints, the range of each value does not actually span the real line). The main simulation parameters are reported in Table II.

At each step t , the observer senses the system by taking a black and white picture of the scene, which is in a space $\mathcal{P} = \{0, \dots, 255\}^{180 \times 360}$. To take the temporal element into account, an observation O_t includes two subsequent pictures, at times $t - 1$ and t , so that the observation space is $\mathcal{O} = \mathcal{P} \times \mathcal{P}$. An example of the transmission process is given in Fig. 2, which shows the original version sensed by the observer (above) and the reconstructed version at the receiver (below) when using a trained VQ-VAE ζ_6 , i.e., an encoder trained with 6 bits per feature, the maximum we consider in this study.

In the CartPole problem, the action space \mathcal{A} contains just two actions *Left* and *Right*, which push the cart to the left or to the right, respectively. At the end of each step, depending on the true state s_t , and on the taken action a_t , the environment will return a deterministic reward

TABLE II: Simulation Parameters.

Parameter	Value	Description
$H \times W$	160×360	Image size
V	7	Number of quantizers
F	8	Number of latent features
D_{emb}	8	Embedding dimension of features
B	256	Batch size
γ	0.95	Discount factor
T	500	Maximum number of steps for an episode
α_{enc}	10^{-3}	VQ-VAE learning rate
α_{Reg}	10^{-4}	Regressor learning rate
α_{A2C}	10^{-4}	A2C learning rate
D	5×10^4	Size of the VQ-VAE training dataset
T_{enc}	100	Encoder training epochs
E_{rob}	2×10^4	Robot policy training episodes
E_{obs}	1×10^5	Observer policy training episodes
E_{test}	1000	Number of test episodes

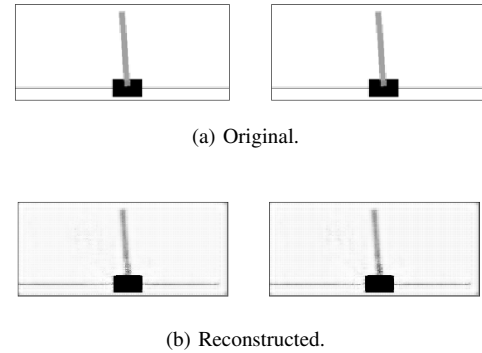


Fig. 2: Example of the original and reconstructed observation.

$R_t = -x_{\text{max}}^{-1}|x_t| - \psi_{\text{max}}^{-1}|\psi_t|$, $\forall t$, where $x_{\text{max}} = 4.8$ m and $\psi_{\text{max}} = \frac{2\pi}{15}$ rad (equivalent to 24°) are the maximum values for the two quantities. If the angle or cart position go outside the boundaries, the episode is over, and the agents do not accumulate any more reward. The goal of the two agents is thus to maximize the cumulative discounted sum of the reward R_t , while limiting the communication cost.

A. The Coding and Decoding Functions

As mentioned in Sec. III-C, the observer can optimize its coding function Λ according to different criteria depending on the considered communication problem. However, as we explained in Sec. IV, optimizing Λ without any parameters is usually not feasible due to the curse of dimensionality on the action space. Consequently, we rely on a pre-trained set \mathcal{V} of VQ-VAE models, whose codebooks are optimized to solve the technical problem, i.e., minimizing the distortion on the observation measured using the MSE: $d_A(o, \hat{o}) = \text{MSE}(o, \hat{o})$. The training performance of the VQ-VAE with ζ_6 , i.e., using the maximum value V of 6 bits per feature, is shown in Fig. 3: the encoder converges to a good reconstruction performance, which can be measured by its *perplexity*. The perplexity is simply $2^{H(p)}$, where $H(p)$ is the entropy of the codeword selection frequency, and a perplexity equal to the number of codewords is the theoretical limit, which is only reached if all codewords are selected with the same probability. The perplexity at convergence is 54.97, which is close to the theoretical limit for a real application.

²https://www.gymnasium.dev/environments/classic_control/cart_pole/

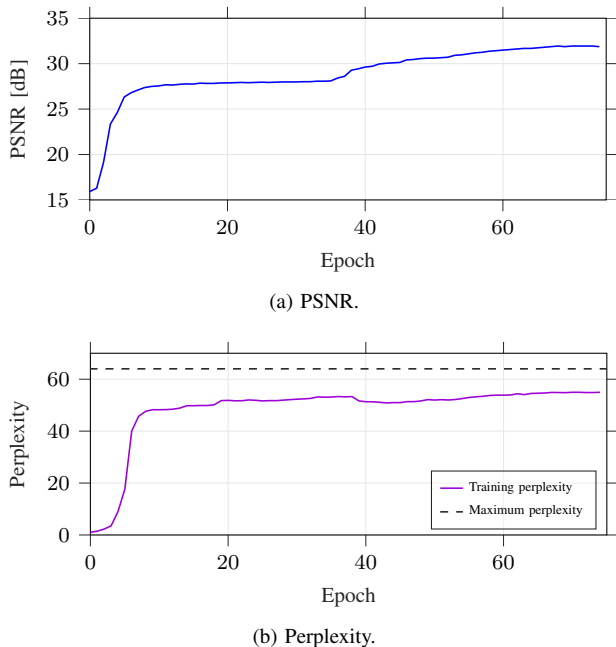


Fig. 3: Training of the VQ-VAE model ζ_6 .

The observer then uses DRL to foresightedly choose the best VQ-VAE ζ_v at each time step, maximizing the expected long-term reward for each communication problem. We train the observer to solve each specific coding problem by designing three different rewards, depending on the considered communication level (A, B, or C):

- 1) *Level A (technical problem)*: The distortion metric for the observer is $d_A(o_t, \hat{o}_t) = -\text{PSNR}(o_t, \hat{o}_t)$, as part of the reward definition from (2). The PSNR is an image fidelity measure proportional to the logarithm of the normalized MSE between the original and reconstructed image;
- 2) *Level B (semantic problem)*: The distortion metric is $d_B(\hat{s}_t^{(o)}, \hat{s}_t^{(r)}) = \text{MSE}(\hat{s}_t^{(o)}, \hat{s}_t^{(r)})$, as part of the reward defined in (3), and the decoder needs to estimate the underlying physical state s_t by minimizing the MSE, i.e., the distance between $\hat{s}_t^{(o)}$ and $\hat{s}_t^{(r)}$ in the semantic space. In our case, the estimator used to obtain the estimates is a pre-trained supervised LSTM neural network;
- 3) *Level C (effective problem)*: In this case, there is no direct distortion metric, and the control performance is used directly as in (4). The policy $\pi^{(r)}$ is given by an actor-critic agent implementing an LSTM architecture, pre-trained using data with the highest available message quality (6 bits per feature).

In this case, the task depends on all the semantic features contained in S_t . However, the 4 components of the state do not carry the same amount of information to the robot: depending on the system conditions, i.e., the state S_t , some pieces of information are more relevant than others.

B. Neural Network Architecture and Training

The VQ-VAE architecture is made with Convolutional Neural Network (CNN) layers to extract latent features and is

TABLE III: Encoder-Decoder parameters.

Layer type	Size	Kernel size	Stride
Encoder			
Conv2d + ReLU	64	10×11	8×9
Conv2d + ReLU	64	12×12	10×10
Conv2d + ReLU	128	3×3	1×1
ResidualStack	2	3×3	1×1
Conv2d	8	3×3	1×1
Decoder			
ResidualStack	2	3×3	1×1
Conv2d + ReLU	128	3×3	1×1
Conv2d + ReLU	64	12×12	10×10
Conv2d	64	10×11	8×9

trained separately before training the control policy. To this end, a dataset of observations is collected through a random policy. We then train an encoding network, the vector quantization layer and the decoder jointly as in the standard VQ-VAE [19]. The first vector quantization layer learned contains the highest number of codewords. Finally, we fix the encoder and the decoder and just train the other vector quantization layers, obtaining multiple quantizers over the same latent space discovered by a common encoder. The hyperparameters used to train the VQ-VAE are reported in Table II. After obtaining the V quantizers, we train the policy using the standard A2C algorithm. Table III shows the Encoder-Decoder layers of the VQ-VAE. In Table IV, the layers of the implemented Regressor and the Actor-critic neural networks are reported. All the NNs are implemented through the Pytorch library. Once the robot policy has been obtained, we can train the observer policy. The observer learns a policy through the same A2C algorithm, but in this case the inputs to the policy are the features before quantization. A unique observer policy is trained for different values of the trade-off parameter β and for different communication levels. For further details on the implementation, training, and testing process, we refer to the publicly available simulation code.³

The additional computational cost of the architecture on the observer side is well within the computational capabilities of even relatively simple embedded devices [32], and even more complex problems can be dealt with by Edge devices. At the same time, training the actor with compressed representations actually reduces its computational burden, as feature extraction is performed by the sender. However, if the sender is significantly computationally constrained, replacing the VQ-VAE with a classical compression scheme such as JPEG might be a good way to reduce the cost and still deliver the benefits of dynamic compression.

In the following, the main measure of the communication cost will be the average message length ℓ_t rather than the index v of the VQ-VAE ζ_v , as we consider it to be a more meaningful physical quantity, particularly when comparing the proposed approach with traditional encoding methods. However, the two parameters are linearly proportional to one another, and since we set $F = 8$, the message length in bytes corresponds to the value of v measured in bits.

³https://github.com/pietro-talli/dynamic_compression_EC

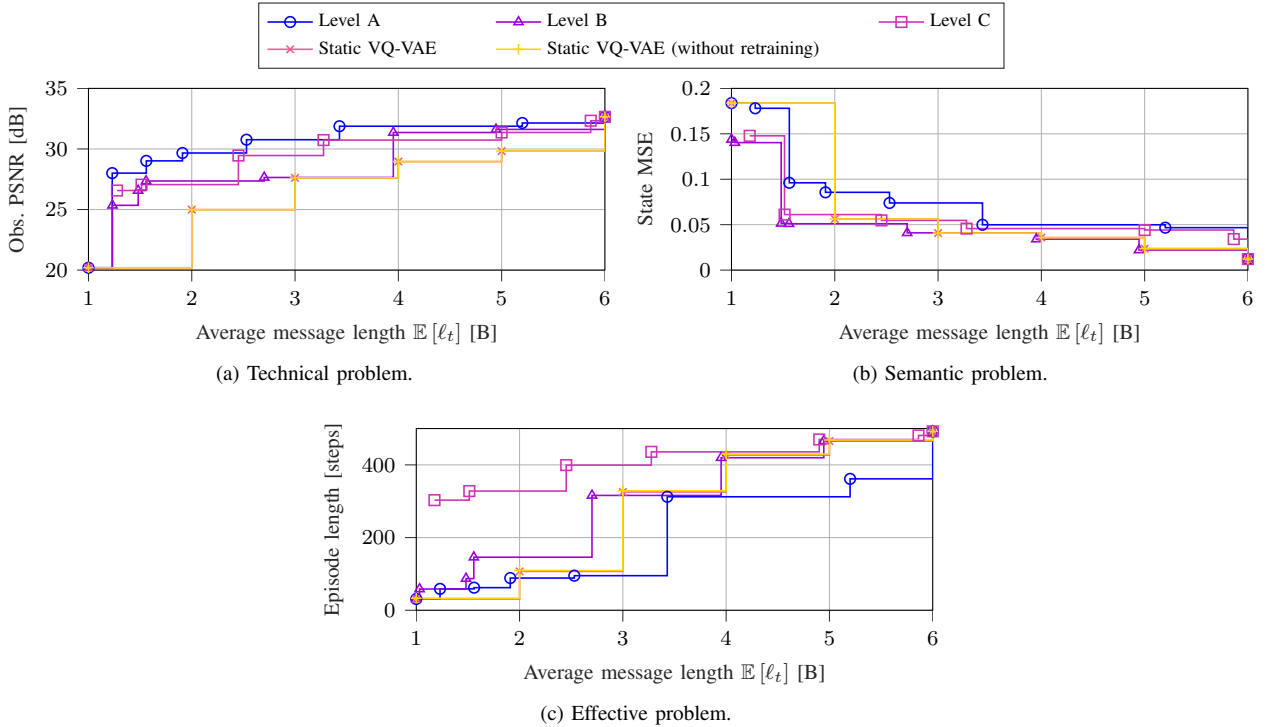


Fig. 4: Performance of the communication schemes on the three levels of the remote POMDP.

TABLE IV: Recurrent architectures.

Layer type	Inputs	Outputs	Description
Regressor			
LSTM + ReLU	64	64	Single recurrent layer
Linear + ReLU	64	128	Hidden layer
Linear	128	1	Output layer
Actor-critic			
LSTM + ReLU	64	64	Single recurrent layer
Linear + ReLU	64	128	Hidden layer
Linear	128	1	Output layer (Value)
Linear + softmax	128	$ A $	Output layer (Policy)

C. Results

We assess the performance of the three different tasks in the CartPole scenario by simulation, measuring the results over 1000 episodes after convergence. Fig. 4 shows the performance of the various schemes over the three problems, compared with a static VQ-VAE with a constant compression level. In the Level C evaluation, we also consider a static VQ-VAE solution in which the robot is not retrained for each v , but is only trained for $v = 6$ bits per feature (i.e., 48 bits per message, as the VQ-VAE considers 8 features) as for the dynamic scheme. We trained the dynamic schemes with different values of the communication cost β , so as to provide a full picture of the adaptation to the trade-off between performance and compression. We also introduce the notion of *Pareto dominance*: an n -dimensional tuple $\eta = (\eta_1, \dots, \eta_n)$ Pareto dominates η' (which we denote as $\eta \succ \eta'$) if and only if:

$$\eta \succ \eta' \iff \exists i : \eta_i > \eta'_i \wedge \eta_j \geq \eta'_j \forall j. \quad (6)$$

We can extend this to schemes with multiple possible configurations. The definition of Pareto dominance for schemes

x and y is: $x \succ y \iff \exists \eta(x) \succ \eta(y) \forall \eta(y)$, i.e., for each configuration of scheme y , there is a setting of x that Pareto dominates it. We can always tune scheme x so that it outperforms any configuration of scheme y on all metrics.

We first consider the technical problem performance, shown in Fig. 4a: as expected, the Level A dynamic compression outperforms all other solutions, and its performance is Pareto dominant with respect to static compression. Interestingly, the Level B and Level C solutions perform worse than static compression: by concentrating on features in the semantic space or the task space, these solutions remove information that could be useful to reconstruct the full observation, but is meaningless for the specified task.

In the semantic problem, shown in Fig. 4b, a lower MSE on the reconstructed state is better, and the Level B solution is Pareto dominant with respect to all others. The Level A solution also Pareto dominates static compression, while the Level C solution only outperforms it for higher compression levels, i.e., on the left side of the graph.

Finally, Fig. 4c shows the performance of the effectiveness problem (Level C), summarized by how long the CartPole system manages to remain within the position and angle limits. The Level C solution significantly outperforms all others, but is not strictly Pareto dominant: when the communication constraint is very tight, setting a static compression and retraining the robot to deal with the specific VQ-VAE used may provide a slight performance advantage. In general, almost perfect control can be achieved with less than half of the average bitrate of the static compressor, which can only reach a similar performance at a much higher communication cost. We also note that, in this case, the Level B solution performs worst: choosing the solution that minimizes the semantic distortion




Fig. 5: Other performance metrics relative to the CartPole control problem.




Fig. 6: Comparison with other methods.

is not always matched to the task, as it considers the state variables as having equal weight, while a higher precision may be required when the quantization error affects the robot's action.

Another analysis is conducted on the way the CartPole is controlled with the different communication policies. Fig. 5 shows the angular Root Mean Squared Deviation (RMSD) (Fig. 5a) and the position RMSD (Fig. 5b), defined as:

$$\text{RMSD}(x) = \sqrt{\frac{1}{t} \sum_{i=1}^t (x_i - x_{\text{target}})^2},$$

where x_{target} is the desired value of the controlled process and x_i is the recorded value of the process at time step i . Both RMSDs are computed with respect to the central and vertical position of the CartPole: $x_{\text{target}} = 0$ and $\psi_{\text{target}} = 0$. These results help to evaluate how well the control dynamics keep the CartPole near the optimal central position and to assess the smoothness of the resulting pole oscillations. It is possible to see that, in general, a higher rate allows to keep the angular RMSD smaller. In particular, in the Level C system, the values are the smallest. However, this comes at the cost of deviating more from the central position, as shown in the figure. The policy prioritizes the stabilization of the pole oscillations, though this requires deviating from the central position. This is because swings in the pole's angle are harder to control due to the instability of the inverted pendulum, and there is a significant risk that the pole might go out of the acceptable range, ending the episode.

D. Comparison with existing compression approaches

In this section, we compare the performance of our proposed approach to that of other methods. Specifically, we show how digital compression techniques, such as JPEG, perform in the same scenario. We also compare other NN-based compression models and evaluate their performance. For the digital compression, we use different sets of parameters combining image resizing, the quality parameter of the JPEG standard, and the number of grayscale color levels. For learning-based compression, we used the CompressAI library, which implements the model proposed in [33].

Fig. 6 shows the performance of other methods with respect to the proposed dynamic feature compression method. It is possible to see that the digital compression scheme does not allow the actor to effectively control the system, as its stability is low, even though the updates are bigger by two orders of magnitude. Obtaining a high control performance would require an extremely high bitrate. On the other hand, the neural compression scheme achieves a higher performance with a limited bitrate, but since the model is a general compression technique designed to compress a wide variety of images, it cannot reach extremely low bitrates. After retraining the scheme on CartPole pictures, it is possible to obtain lower bitrates while improving the resulting control performance. Our approach, which is directly trained on the CartPole task, outperforms all others; however, we do not claim that VQ-VAE is the best compression technique for all scenarios. The main contribution of this work is not in the specific compression scheme, but rather in the dynamic and goal-oriented adaptation of the compression parameters for each transmitted update, which could be directly applied to different NN architectures and even JPEG.

E. Analysis of the communication policy

We can then use an explainability approach to gain further insight into how effective communication operates. Fig. 7 shows the distribution of the quantization level selected by an observer trained for each communication Level (A, B, and C). We note that we adapted the scale of β for the three levels, so as to obtain comparable results: as the reward process takes values in different ranges (e.g., the PSNR is in dB while the reward of the MDP is between -1 and 1), using the same transmission cost would result in very different outcomes.

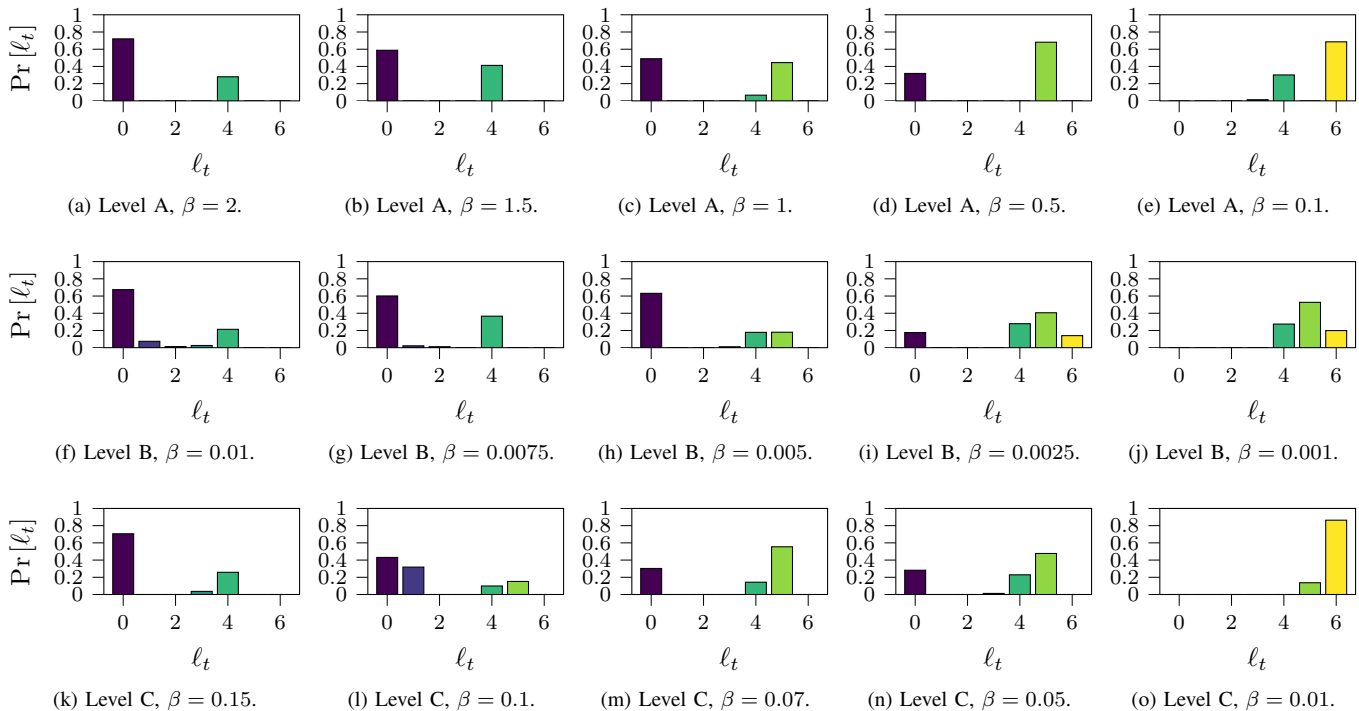


Fig. 7: Distribution of the selected compression levels.

We then chose to rescale the transmission cost parameter to have the full range of average bitrates at each communication level. The similarity in the compression level distributions at the three levels is striking. For lower values of β , the observer selects ζ_{\emptyset} , which corresponds to no transmission, more often. As β decreases, the communication cost becomes lower, and thus the observer chooses longer messages more often. Another common pattern is that quantizing features using 1, 2 or 3 bits, i.e., sending messages of 1, 2 or 3 bytes, respectively, is a rare choice. This shows that the memory implemented implicitly in the system through the LSTM is powerful enough to obtain adequate beliefs based on past messages, so that the observer can rely on it and not send anything. A roughly quantized update has a relatively low value, as its novelty is limited, and transmitting intermittent updates at a higher quality results in a better performance.

However, the real difference between the three policies is given by *when* they decide not to transmit. Therefore, we propose an analysis based on the visualization of the observer policy and the receiver policy. In Fig. 8, four colormaps show different policies projected in the same domain: the pole angle on the x-axis and the cart velocity on the y-axis. More specifically, we quantize the projected state into cells and show the policy of the robot and the observer in each cell. Fig. 8a shows the robot actions, averaged among 10^6 samples. Since actions in the CartPole problem are binary, we represent the probability of choosing action Right in a range between 0 and 1. Fig. 8b depicts the entropy of the robot policy. We consider the action probability in the previous figure and compute the action entropy as follows:

$$H(a) = -p(0) \log_2(p(0)) - p(1) \log_2(p(1)),$$

where $p(a)$ is empirically estimated by counting the number of times each action is chosen when the state is in the projected cell. Fig. 8c and Fig. 8d show the average length of the update packets, i.e., the number of Bytes transmitted in each cell when optimizing for Levels A and C, respectively. This can be seen as the average number of bits allocated for each projected slice of the state space. Fig. 9 shows the same results, but for a different physical state projection, mapping the angle ψ on the x-axis and the pole angular velocity $\dot{\psi}$ on the y-axis.

In both figures, there is a strong correspondence between the states where the robot entropy is higher and the states where the Level C policy allocates a higher number of bits. This confirms that an effective observer policy manages to discriminate the uncertainty at the robot side. In regions of the state space where it is more difficult to retrieve the correct action, i.e., the action entropy is higher, the observer will provide the robot with more precise information by sending longer messages. There are regions where the robot action is always the same, e.g., whenever the cart is moving fast and the tip of the pole is pointing to the same side the cart is moving towards. In these cases, the entropy is extremely low and the transmitter can avoid sending new updates to the robot. This is due to the fact that, even if the estimated state at the receiver differs from the observed one, the action to perform remains the same and will be to push further the cart to try to get the pole more vertical. Recalling (5), we note that if $Q(\xi_t^{(r)}, \pi^{(r)}(\xi_t^{(r)}))$ is very sensitive to small variations in $\xi_t^{(r)}$, then the gap in (5) is going to be significant, leading the observer to choose to send precise information. In principle, a Level C transmitter could reduce the message length or even avoid transmission as long as the robot is able to choose the

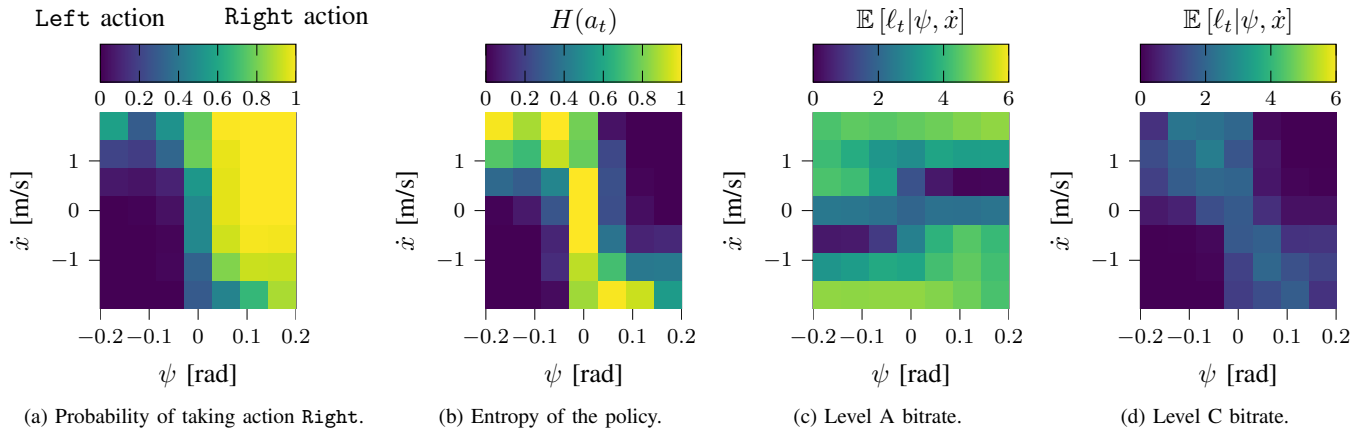


Fig. 8: Analysis of the transmission policy as a function of the pole angle ψ and cart linear velocity \dot{x} . The bitrate is measured in Bytes per transmission.

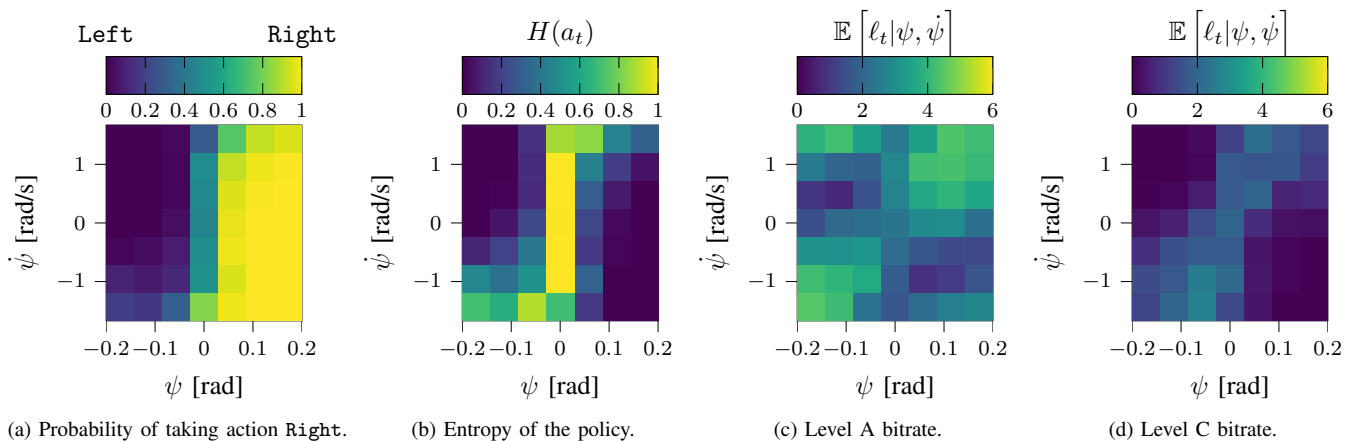


Fig. 9: Analysis of the transmission policy as a function of the pole angle ψ and angular velocity $\dot{\psi}$. The bitrate is measured in Bytes per transmission.

correct actions, even though its belief is incorrect. An optimal communication scheme approximately follows

$$\ell_t \propto V\left(\xi_t^{\text{pri},(r)}, m\right),$$

which means that the message length is roughly proportional to VoI. This concept might be used when defining a heuristic policy, which behaves similarly to the effective communication policy, but is much simpler to design and implement. Note that this condition includes two separate cases in which a Level C observer chooses not to transmit, while Level A and B transmitters would send precise data:

- The action corresponding to the prior belief is the same as the one after the updating message. In this case, the VoI of the communicated message is low and thus we can lower ℓ_t ;
- The action is different after the communicated message, but the long-term rewards are close enough that the robot is not going to benefit too much from choosing the other action. Even in this case, sending less information is not going to affect the control performance significantly.

These cases cannot be taken into account at Levels A and B. Indeed, the Level A policy shown in Fig. 8c tends to allocate communication resources in the states where the

picture is changing more rapidly, so that the memory available to the robot is less useful to estimate the current observation, regardless of the correct action. As the cart speed \dot{x} increases along the y-axis, the number of bits also increases. The same reasoning can be applied to the results in Fig. 9.

Another general principle that we can deduce for an effective policy is that it should be aware of variations of the value function with respect to the belief. If the value function is strongly affected by small perturbations of the belief, then the effective policy should communicate more information in order to reduce the discrepancy between $\xi_t^{(o)}$ and $\xi_t^{(r)}$. This reasoning can be intuitively understood by looking at the differential of the robot's value function $Q\left(\xi_t^{(r)}, \pi^{(r)}\left(\xi_t^{(r)}\right)\right)$ with respect to changes in its belief distribution $\xi_t^{(r)}$. When this value is large, an inaccurate estimation of the state would cause a poor estimation of the value function, which may in turn cause the robot to choose a low-quality action.

In Fig. 10, we provide an analysis of the communication strategy with respect to different AoI values. This allows us to show how the memory of the robot and of the observer plays a crucial role on the communication decisions. In particular, we consider five values of the AoI: AoI = 0 indicates that a message of any length was transmitted in the previous time

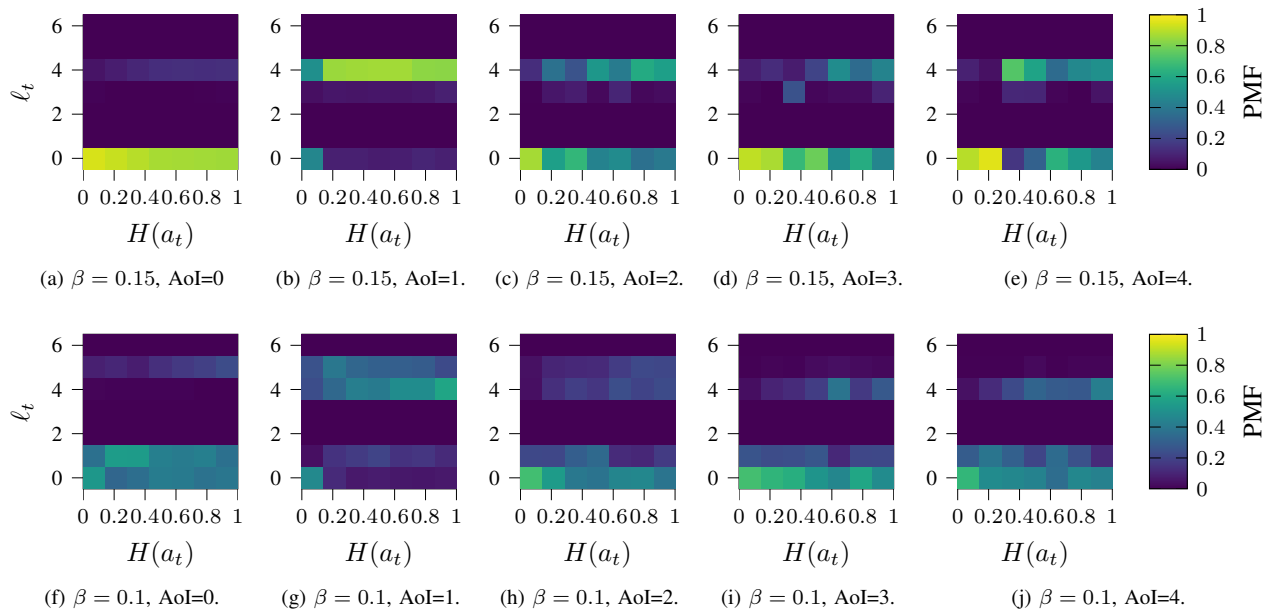


Fig. 10: Level C observer action distribution for different robot action entropy levels and values of β . The color of each cell represents the empirical probability of choosing a packet length ℓ_t for a given entropy $H(a_t)$ and AoI. Each column of each subfigure then represents a conditional Probability Mass Function (PMF).

step. AoI = Δ with $\Delta \in \{1, 2, 3, 4\}$ means that no messages have been received by the observer for Δ time steps since the last received message. This is a measure of how up to date the memory of the robot is, allowing us to evaluate the next choice of the observer for a given age. We then consider the distribution of the observer actions (y-axis) with respect to different ranges of the robot action entropy (x-axis). This means that, for each entropy interval, we count the number of times each action is performed, in order to obtain an empirical distribution. The columns are normalized so that each cell shows the probability that the observer chooses a specific ℓ_t whenever the robot action entropy falls within the corresponding interval, for different values of the AoI.

Fig. 10a clearly shows that, if there was a transmission in the previous time step (AoI = 0), it is very unlikely that the system is going to be updated again in the current time step. As remarked above, Figs. 10a-e show the case with $\beta = 0.15$, for which the observer almost always picks ζ_4 , sending 4-byte messages. On the other hand, Figs. 10f-j show the case with $\beta = 0.1$, in which the agent sometimes selects other codebooks due to the lower transmission costs. The observer often chooses to communicate if AoI = 1, with an exception if the system is in a very low entropy state, in which case the probability of communicating using ζ_4 is similar to the one corresponding to action ζ_\emptyset . If we look at the behavior for higher values of the AoI, we can notice a general trend: communication is more likely to happen in higher entropy states than in lower entropy ones. This shows that the observer policy understands the cases where the state has to be precisely estimated by the robot to choose its action correctly. Additionally, the probability of transmitting an update actually *decreases* when the AoI reaches higher values. The observer will skip several consecutive transmission opportunities only

in two cases: either the system state is highly predictable, and the actor can rely on its past knowledge to get a precise estimate of it, or the two actions are almost equivalent, e.g., if the pole is balanced vertically. The former case is more likely to be a low-entropy state, while the latter is high-entropy, but has a small difference between the rewards for the two actions. We can see that the trend holds for different values of β by looking at Figs. 10f-j: if we decrease the value of β , the observer tends to transmit more often, and use higher message lengths when it transmits, but the general tendency to transmit more whenever the robot action entropy is high clearly holds. This final analysis allows us to get an easy heuristic for effective communication when the value function is not available. This behavior is observed for each subfigure individually, with some exceptions, but not across the subfigures, as the relative transmission probability also depends on the age of the information available to the receiver, which could affect the quality of its estimate. We also note that there are other factors that affect transmitter decisions, such as the relative difference between the q-values of the two actions: the VoI of an update might be low if the two actions are almost equivalent, even if the entropy of the receiver policy is high.

Fig. 11 shows that the Level A policy allocates communication resources without considering the entropy of the control actions. While Fig. 10 showed a clear monotonic trend in the probability of selecting ζ_4 , which increased as the entropy increased, the pattern is much weaker in this case, with the exception of Fig. 11a and Fig. 11c, which show a monotonic behavior similar to that of the Level C policy. The value of β for this was chosen to get a similar overall bitrate (and, as we discussed, a similar overall action distribution) to the Level C case with $\beta = 0.15$. As we discussed above, there is a weak correlation between the action entropy and




Fig. 11: Level A observer action distribution for different robot action entropy levels with $\beta = 1$. The color of each cell represents the empirical probability of choosing a packet length ℓ_t for a given entropy $H(a_t)$ and AoI. Each column of each subfigure then represents a conditional PMF.

features such as the angular and cart velocities, but it is the latter that the Level A policy considers: as the difficulty of accurately reconstructing the image increases with the speed of the CartPole system, more unstable states have more frequent transmissions. Fig. 11 shows that the Level A system exhibits an almost binary threshold-based behavior: the sensor either selects $\zeta_v = 5$ or does not transmit at all. Conversely, Level C yields a similar average communication rate, but does so by transmitting shorter messages more frequently. A proper understanding of these behaviors and its generalization to other contexts requires further study and is outside the main scope of this work. However, the logarithmic shape of the Level A reward function (PSNR), combined with a linear cost function of the message length, may tend to favor the transmission of fewer, more accurate messages. On the other hand, Level C's reward function requires the observer to just transmit enough information for the robot to pick the correct action, leading to shorter, more frequent messages.

VI. CONCLUSION

In this work, we presented a dynamic feature compression scheme that can exploit an ensemble VQ-VAE to solve the semantic and effective communication problems. The dynamic scheme outperforms fixed quantization and can be trained automatically with limited feedback, unlike emergent communication models that are unable to deal with complex tasks. The choices made by the observer are clearly tied to the control policy of the robot it aims to help, significantly outperforming a simpler optimization that does not take into account the semantic and effective problems. We also analyzed the optimal policies to draw insights on their decisions, showing that the Level C optimization indeed considers the robot's policy.

A natural extension of this model is to consider more complex tasks and wider communication channels, corresponding to realistic control scenarios, or scenarios with multiple transmitters with partial information about each other and the robot. Considering a more realistic channel model, which has a loss probability and time-varying statistics in addition to the transmission cost, would also be an interesting direction, joining our model with Joint Source Channel Coding (JSCC) theory. Dynamically adapting JSCC parameters with the goal of helping a remote DRL agent would be a natural extension of our proposed approach for more realistic wireless scenarios.

Another interesting direction for future work is to consider joint training of the robot and the observer, or cases with partial information available at both transmitter and receiver, as well as considering more extensive theoretical performance guarantees in relatively simple and tractable systems.

REFERENCES

- [1] C. E. Shannon and W. Weaver, *The mathematical theory of communication*. University of Illinois Press, Sep. 1949.
- [2] D. Gündüz, Z. Qin, I. E. Aguerri, H. S. Dhillon, Z. Yang, A. Yener, K. K. Wong, and C.-B. Chae, "Beyond transmitting bits: Context, semantics, and task-oriented communications," *IEEE Journal on Selected Areas in Communications*, vol. 41, pp. 5–41, Jan. 2023.
- [3] P. Popovski, O. Simeone, F. Boccardi, D. Gündüz, and O. Sahin, "Semantic-effectiveness filtering and control for post-5G wireless connectivity," *Journal of the Indian Institute of Science*, vol. 100, no. 2, pp. 435–443, Apr. 2020.
- [4] P. A. Stavrou and M. Kountouris, "The role of fidelity in goal-oriented semantic communication: A rate distortion approach," *IEEE Transactions on Communications*, vol. 71, no. 7, pp. 3918–3931, Jul. 2023.
- [5] S. Wan, Z. Gu, and Q. Ni, "Cognitive computing and wireless communications on the edge for healthcare service robots," *Computer Communications*, vol. 149, pp. 99–106, Jan. 2020.
- [6] R. D. Yates, Y. Sun, D. Richard Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of Information: An Introduction and Survey," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1183–1210, Mar. 2021.
- [7] Y. Shao, Q. Cao, and D. Gündüz, "A Theory of Semantic Communication," *arXiv:2212.01485*, Dec. 2022.
- [8] E. Beck, C. Bockelmann, and A. Dekorsy, "Semantic communication: An information bottleneck view," *arXiv:2204.13366*, Apr. 2022.
- [9] J. Shao, Y. Mao, and J. Zhang, "Learning task-oriented communication for edge inference: An information bottleneck approach," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 1, pp. 197–211, Jan. 2022.
- [10] E. Uysal, O. Kaya, A. Ephremides, J. Gross, M. Codreanu, P. Popovski, M. Assad, G. Liva, A. Munari, B. Soret *et al.*, "Semantic communications in networked systems: A data significance perspective," *IEEE Network*, vol. 36, no. 4, pp. 233–240, Oct. 2022.
- [11] F. Pase, S. Kobus, D. Gündüz, and M. Zorzi, "Semantic communication of learnable concepts," *International Symposium on Information Theory (ISIT)*, Jun. 2023.
- [12] T.-Y. Tung, S. Kobus, J. P. Roig, and 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*, vol. 39, no. 8, pp. 2590–2603, Jun. 2021.
- [13] D. Kim, S. Moon, D. Hostallero, W. J. Kang, T. Lee, K. Son, and Y. Yi, "Learning to schedule communication in multi-agent reinforcement learning," in *7th International Conference on Learning Representations (ICLR)*, May 2019.
- [14] X. Zheng, S. Zhou, and Z. Niu, "Urgency of Information for Context-Aware Timely Status Updates in Remote Control Systems," *IEEE Transactions on Wireless Communications*, vol. 19, no. 11, pp. 7237–7250, Jul. 2020.

- [15] J. N. Foerster, Y. M. Assael, N. de Freitas, and S. Whiteson, "Learning to communicate with deep multi-agent reinforcement learning," in *30th International Conference on Neural Information Processing Systems (NeurIPS)*, Dec. 2016.
- [16] A. Goyal, R. Islam, D. Strouse, Z. Ahmed, H. Larochelle, M. Botvinick, S. Levine, and Y. Bengio, "Transfer and exploration via the information bottleneck," in *7th International Conference on Learning Representations (ICLR)*, May 2019.
- [17] F. Pase, D. Gündüz, and M. Zorzi, "Rate-constrained remote contextual bandits," *IEEE J. Sel. Areas Inform. Theory*, vol. 3, no. 4, pp. 789–802, Dec. 2022.
- [18] P. A. Stavrou, J. Østergaard, and C. D. Charalambous, "Zero-delay rate distortion via filtering for vector-valued gaussian sources," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 5, pp. 841–856, Oct. 2018.
- [19] A. Van Den Oord, O. Vinyals *et al.*, "Neural discrete representation learning," in *31st International Conference on Neural Information Processing Systems (NeurIPS)*, Dec. 2017.
- [20] P. Talli, F. Pase, F. Chiariotti, A. Zanella, and M. Zorzi, "Semantic and effective communication for remote control tasks with dynamic feature compression," in *16th Int. Worksh. Wireless Sensing Actuating Rob. Networks (INFOCOM WiSARN)*. IEEE, May 2023.
- [21] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "The age of information in a discrete time queue: Stationary distribution and non-linear age mean analysis," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1352–1364, May 2021.
- [22] Z. Wang, M.-A. Badiu, and J. P. Coon, "A framework for characterizing the value of information in hidden Markov models," *IEEE Transactions on Information Theory*, vol. 68, no. 8, pp. 5203–5216, Aug. 2022.
- [23] E. Fountoulakis, N. Pappas, and M. Kountouris, "Goal-oriented policies for cost of actuation error minimization in wireless autonomous systems," *IEEE Commun. Lett. (Early Access)*, Jun. 2023.
- [24] E. Boursoulatzé, D. Burth Kurka, and D. Gündüz, "Deep joint source-channel coding for wireless image transmission," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 3, pp. 567–579, May 2019.
- [25] H. Xie, Z. Qin, G. Y. Li, and B.-H. Juang, "Deep learning enabled semantic communication systems," *IEEE Transactions on Signal Processing*, vol. 69, pp. 2663–2675, Apr. 2021.
- [26] F. Mason, F. Chiariotti, A. Zanella, and P. Popovski, "Multi-agent reinforcement learning for coordinating communication and control," *IEEE Transactions on Cognitive Communications and Networking*, Apr. 2024.
- [27] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, no. 1, pp. 99–134, May 1998.
- [28] E. J. Sondik, "The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs," *Oper. Res.*, vol. 26, no. 2, pp. 282–304, Apr. 1978.
- [29] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, Nov. 2018.
- [30] T. M. Cover and J. A. Thomas, *Elements of Information Theory, Second Edition*. Wiley-Interscience, Sep. 2006.
- [31] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *2nd International Conference on Learning Representations (ICLR)*, May 2014.
- [32] P. Kang and J. Jo, "Benchmarking modern edge devices for AI applications," *IEICE Transactions on Information and Systems*, vol. 104, no. 3, pp. 394–403, Mar. 2021.
- [33] Z. Cheng, H. Sun, M. Takeuchi, and J. Katto, "Learned image compression with discretized Gaussian mixture likelihoods and attention modules," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Jun. 2020.
- [34] V. Mnih, K. Kavukcuoglu, D. Silver *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015.

APPENDIX

THE INFORMATION BOTTLENECK VIEW

We can also consider another perspective on the observer's choices, using information bottleneck theory. We define a *sufficient statistic* $i(s)$ of any given state $s \in \mathcal{S}$, which is enough to determine the robot's performance in that state.

Denoting the number of bits required to represent a realization of random variable X as $b(X)$, we consider a case in which:

$$b(i(S)) < b(S) < b(O).$$

Indeed, the observation may contain much more information than needed to estimate the state [30], and lossily compressing the message to preserve the relevant information, removing redundant or irrelevant details, can ease communication requirements without any performance loss. We can also observe that $i(S) \rightarrow S \rightarrow O$ is a Markov chain. The random quantity $i(S)$ represents the minimal description of the system with respect to the robot's task, i.e., no additional data computed from S adds meaningful information for the robot's policy. The state S may also include task-irrelevant physical information on the system. However, both S and $i(S)$ are unknown quantities, as the observer only receives a noisy and high-dimensional representation of S through O . This is a well-known issue in DRL: in the original paper presenting the Deep Q-Network (DQN) architecture [34], the agent could only observe the screen while playing classic arcade videogames, and did not have access to the much more compact internal state representation of the game. Introducing communication and dynamic encoding adds another layer of complexity.

We can then consider the case in which communication is limited to a maximum length of L bits, i.e., to $2^{L+1} - 1$ messages, considering all possible lengths lower than or equal to L , including no communication. Naturally, this assumes that the receiver has a way to discriminate between messages of different lengths, e.g., through a MAC layer header. The channel is ideal, i.e., instantaneous and error-free, but it includes a constant cost per bit β , as in the observer reward we gave in the previous section. Consequently, the problem introduces an information bottleneck between the observation O_t and the estimate \hat{o}_t that the robot can make, based on the message M_t conveyed through the channel. If we define a distortion measure over the observation space $d_A : \mathcal{O} \rightarrow \mathbb{R}^+$, any communication introduces a non-zero distortion $d_A(o, \hat{o})$ whenever $b(o) > L$, whose theoretical asymptotic limits are given by rate-distortion theory [30]. If we also consider *memory*, i.e., the use of past messages in the estimation of \hat{o} , the mutual information between o and the previous messages can be used to reduce the distortion, improving the estimate.

In the semantic problem, the aim is to extrapolate the real physical state of the system S_t from the compressed observation M_t , which can be a complex stochastic function. In general, the real state lies in a low-dimensional semantic space \mathcal{S} . The term "semantic" is motivated by fact that, in this case, the observer is not just transmitting pure sensory data, but rather some meaningful piece of physical information about the system. Consequently, the distortion to be considered in this case can be represented by a measure $d_B : \mathcal{S}^2 \rightarrow \mathbb{R}$ over the semantic space, so that the distortion $d_B(\hat{s}_t^{(o)}, \hat{s}_t^{(r)})$ is computed between the observer's best estimate of the state and the one performed by the robot based on M_t , and on its memory of past messages. Finally, to be even more efficient and specific with respect to the task, the observer may optimize the message M_t to minimize a distortion measure $d_C(i(\xi_t^{(o)}), i(\xi_t^{(r)}))$ between the effective representation

of the observer's belief on the state, which contains only the task-specific information, and the knowledge available to the robot. Naturally, any message instance $m_t \in \mathcal{M}$ must be at most L bits long in order to respect the constraint. However, defining the sufficient statistic $i(\xi_t^{(o)})$ may be highly complex and problem-dependent, and using the robot's reward as a direct performance measure is significantly more direct, with the same guarantees.



Michele Zorzi (F'07) received his Ph.D. in electrical engineering from the University of Padova, Italy, in 1994. After spending three years with the Center for Wireless Communications at UCSD, in 1998 he joined the School of Engineering of the University of Ferrara, Italy, where he became a professor in 2000. Since November 2003 he has been at the Information Engineering Department at the University of Padova. He was the Editor-in-Chief of several IEEE journals, and has served the IEEE Communications Society in various leadership roles.



Pietro Talli [S'23] is a PhD student at the Department of Information Engineering (DEI), University of Padova (Italy). He received his Bachelor's Degree in Information Engineering in 2020 and the Master's Degree in ICT for Internet and Multimedia in 2022. He started his PhD in 2023 working on Semantic and Goal-oriented Communications. His research interests are on machine learning for communication and multi-agent systems.



Francesco Pase [GS'20] is currently the Lead AI Research Engineer at NEWTWEN and received his PhD from the University of Padova, Italy, in 2024. He spent the summer 2020 working as a Research Intern at InstaDeep Ltd, London, UK, conducting research on deep reinforcement learning. During the PhD, he spent 6 months as a visiting researcher at the Imperial College London in 2022, and 3 months at Nokia Bell Labs, Cambridge, UK, in 2023. His research interests are in the fields of information theory, distributed learning, wireless communications, and physics-informed machine learning.



Federico Chiariotti [S'15-M'19] is currently an Assistant Professor in the Department of Information Engineering, University of Padova, Italy, where he also received his PhD in information engineering in 2019. From 2020 to 2022, he worked on semantic communication and Age of Information at Aalborg University, Denmark. He has published more than 80 peer-reviewed papers and received the Best Paper Award in 4 conferences. His research focuses on Age of Information, the use of machine learning in networks, and on goal-oriented communications.



Andrea Zanella [S'98-M'01-SM'13] is a Full Professor at the Department of Information Engineering (DEI), University of Padova (Italy). He received the Laurea degree in computer engineering in 1998 and the PhD in 2001 from the same University. He is one of the coordinators of the SIGnals and NETworking (SIGNET) research lab. His long-established research activities are in the fields of protocol design, optimization, and performance evaluation of wired and wireless networks.