



HAL
open science

Pragmatic Goal-Oriented Communications Under Semantic-Effectiveness Channel Errors

Tomás Huttebraucker, Mohamed Sana, Emilio Calvanese Strinati

► **To cite this version:**

Tomás Huttebraucker, Mohamed Sana, Emilio Calvanese Strinati. Pragmatic Goal-Oriented Communications Under Semantic-Effectiveness Channel Errors. 2024 IEEE 21st Consumer Communications & Networking Conference (CCNC), Jan 2024, Las Vegas, United States. pp.683-689, 10.1109/CCNC51664.2024.10454728 . cea-04549375

HAL Id: cea-04549375

<https://cea.hal.science/cea-04549375>

Submitted on 30 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pragmatic Goal-Oriented Communications under Semantic-Effectiveness Channel Errors

Tomás Huttebraucker, Mohamed Sana, Emilio Calvanese Strinati
CEA-Leti, Université Grenoble Alpes, F-38000 Grenoble, France
Email : {tomas.huttebraucker; mohamed.sana; emilio.calvanese-strinati}@cea.fr

Abstract—In forthcoming AI-assisted 6G networks, integrating semantic, pragmatic, and goal-oriented communication strategies becomes imperative. This integration will enable sensing, transmission, and processing of exclusively pertinent task data, ensuring conveyed information possesses understandable, pragmatic semantic significance, aligning with destination needs and goals. Without doubt, no communication is error free. Within this context, besides errors stemming from typical wireless communication dynamics, potential distortions between transmitter-intended and receiver-interpreted meanings can emerge due to limitations in semantic processing capabilities, as well as language and knowledge representation disparities between transmitters and receivers. The main contribution of this paper is two-fold. First, it proposes and details a novel mathematical modeling of errors stemming from language mismatches at both semantic and effectiveness levels. Second, it provides a novel algorithmic solution to counteract these types of errors which leverages optimal transport theory. Our numerical results show the potential of the proposed mechanism to compensate for language mismatches, thereby enhancing the attainability of reliable communication under noisy communication environments.

I. INTRODUCTION

AI and beyond-5G communication systems need each other to grow but still follow separate paths. In AI algorithms, it is often assumed that communications are ideal, latency-power-error free, trustworthy, and always available (on demand). This is critically far from operational reality. On the other hand, 5G is not AI-native [1][2]: Communication is designed to send raw data rather than AI understandable knowledge. While data is decisive to feed reasoning engines, being gathered-transferred-processed from scattered sources to update models, to accumulate knowledge and attain goals, communications are not engineered to be goal-oriented. The current approach is to encode data into symbols, ensuring symbol level accurate recovery even at ultra-high data rate exchange, regardless of their conveyed meaning [3] and agnostic to what cooperating reasoning agents need and can understand. Conversely, an AI-native approach requires to convey meaning or to accomplish a goal with a reasoning engine. Thus, what really matters is the impact that the received symbols have on the interpretation of the meaning intended by the transmitter and/or on the accomplishment of a common goal. In this context a new impulse in research has started: *semantic* and *goal-oriented* communications have been proposed for future AI-native 6G networks [3] [4],[5], based on the seminal works of Shannon and Weaver [6] and Bar-Hillel and R. Carnap [7], then on the more recent work of J. Bao [8]. This new area of research

focuses on developing theories and algorithmic solutions for extracting, sharing and interpreting distilled information representing the significance of data, thus conveying the essential meaning to enable effective task execution. Such exchange of semantics that allows cooperation between intelligent agents requires a shared language which serves as a tool to endow real world observations with meaning that can be interpreted and thus enables decision making. The design of such language is a challenging task and it often relies on Machine Learning (ML) algorithms. ML-based language design has been extensively studied [9], [10]. Research indicates that incorporating communication as part of the learning procedure results in superior performance compared to human-designed communication protocols [11] [12].

Most literature on the state of the art on semantic and goal oriented communication assumes that interacting agents share a common language. However, this might be an unrealistic assumption. Indeed, in future 6G communication multiple heterogeneous autonomous agents will communicate and collaborate, and they may utilize distinct languages. When agents use different languages, *semantic noise* causes interpretation errors and performance degradation [13]. Similarly to [14], this work addresses the problem of communication errors due to language mismatch between a transmitter and a receiver. We model the communication channel in three levels, the syntactic (wireless) level, the semantic (meaning) level and the effectiveness (goal) level. A language deals with all three levels of the communication channel: it has to enable the completion of the task (*effectiveness*), it has to allow for meaning transmission (*semantic*) and it has to account for possible channel noise (*syntactic*). In order to study heterogeneous language interactions, we treat a language mismatch as a semantic level problem and explore its consequences at the effectiveness level.

The **contribution of the paper** is three-fold: (1) We extend the work from [14] to language mismatch between agents in a Markov Decision Process (MDP) context, (2) we propose a novel mathematical model for errors stemming from language mismatch at both semantic and effectiveness levels and, (3) we propose an equalization algorithm that leverages optimal transport theory to counteract language mismatches and show its efficacy on a practical example.

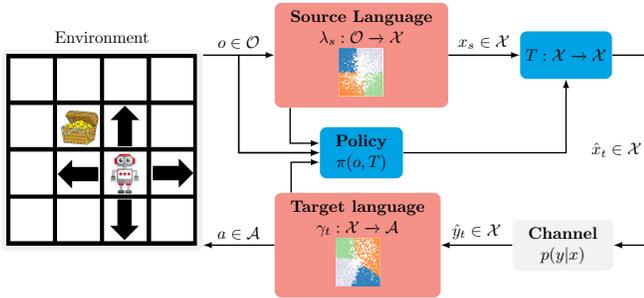


Figure 1. System model.

II. SYSTEM MODEL

We will consider the system model shown in Fig. 1. A transmitter observes a world, here the states of a stochastic environment, and encodes the underlying information for a distant receiver, which rules it to an action to complete a task. Communication takes place over an error-prone wireless channel, in several rounds, until the task is completed or it fails. To achieve the task, both encoder and decoder employ distinct languages (different logic), here artificially modeled with Neural Network (NN), which we respectively refer to as *language generator* and *language interpreter*.

Language generator: A language generator λ maps an observation of the world into semantic representation: a latent representation capturing the intrinsic structure of the data, which is useful for solving downstream tasks. We define such mapping with $\lambda : \mathcal{O} \rightarrow \mathcal{X}$, where \mathcal{O} denotes the observation space and \mathcal{X} the semantic space. Both \mathcal{O} and \mathcal{X} are topological spaces, we will assume \mathcal{O} to be discrete and \mathcal{X} to be continuous, however, our results are easily extended to different spaces. Thus, given an observation $o \in \mathcal{O}$, the language generator outputs the associated semantic representation $x \in \mathcal{X}$, later referred to as a semantic symbol. It is worth noting that λ can, but is not required to be a one-to-one mapping. It can be a many-to-one or one-to-many mapping. In addition, λ is not necessarily deterministic. It can be stochastic, in which case, $\lambda(o) = x \sim \mu_\lambda(x'|o)$. Generally speaking, we can think of a deterministic mapping as a particular case of a stochastic one where $\mu_\lambda(x|o) = \delta(x - \lambda(o))$. Hence, given a distribution μ over the observation space \mathcal{O} , the language generator λ defines a distribution over the semantic space \mathcal{X} as:

$$\mu_\lambda(x) = \sum_{o \in \mathcal{O}} \mu_\lambda(x|o)\mu(o). \quad (1)$$

A language generator defines a way to encode information from the observation space into the semantic space. This effectively creates a partition of the semantic space. A partition of the semantic space $P = \{P_1, P_2, \dots, P_N\}$ is a collection of measurable sets that cover \mathcal{X} , where the measure is μ_λ . The set $P_i \in \mathcal{X}$ is referred to as the i -th atom of the partition and $\mu_\lambda(P_i)$ is its associated probability measure¹. Each atom

¹Here, $\mu_\lambda(P_i)$ can be interpreted as the probability that given a random observation o , $x = \lambda(o)$ belongs to P_i .

P_i is associated to a given semantic meaning. For example, in an image classification task, each atom is related to a different visual feature present of the images, meaning that if an image has a certain visual feature, its semantic representation will reside in the corresponding atom. In general, the semantic space partition depends only on the language generator and two different generators can partition the language differently for the same task. However, it has been observed that task specific optimal information encoding strategies usually share common characteristics [15].

Language interpreter: A language interpreter γ transforms a semantic representation into a task specific action in each communication round. More specifically, a language interpreter defines a strategy for interpreting the semantic partition defined by the language generator to solve the downstream task. We define such selection strategy with $\gamma : \mathcal{X} \rightarrow \mathcal{A}$, where \mathcal{A} denotes the action space, which we will assume to be discrete. Thus, given a received semantic symbol $x \in \mathcal{X}$, the language interpreter selects an action $a = \gamma(x) \in \mathcal{A}$, possibly sampled from a probability distribution $\mu_\gamma(a'|x)$. In the image classification task example used for the language generator, the language interpreter leverages the semantic atom features of the generator's partition to infer a specific class. Hereafter, we denote a language (or a communication strategy) as a tuple $\ell = (\lambda, \gamma, \mu, \mathcal{O}, \mathcal{X}, \mathcal{A})$.

A. Emergence of language and communication effectiveness

An effective communication between the encoder and decoder requires a joint learning procedure, during which a *shared* language emerges to solve the downstream task. An efficient language generator should define a minimal representation, transparent and sufficient that enables the language interpreter to accomplish the task. A representation is sufficient if it encodes all the information relevant for the decision making and transparent if the encoding process is understandable by the interpreter, i.e. the encoding and decoding logic is the same. Minimality is related to the compression capabilities of the encoder and is directly linked to a particular task: a representation is minimal if it encodes only the task relevant information.

Ideally, the learned languages logic should only be affected by the task definition, the designed loss function and other architecture-specific constraints such as the characteristics of the semantic space \mathcal{X} . However, multiple other factors may affect the outcomes of the training such as the random weights initialization, training hyperparameters, or other sources of randomness including syntactic channel noise [15]. As a result, even with the same NN architecture and constraints, different training processes can yield different languages which nevertheless achieve similar performances on the task. When the language generator and interpreter are the result of different training procedures, communication may suffer from language mismatch or incompatibility that stimulates effectiveness errors. This is the focus of our work, to define a novel framework to counteract and limit the critical errors in interpretation to achieve effective communication.

B. Semantic channel and effectiveness channel

In this section, we propose to model the semantic and effectiveness channel between an encoder-decoder pair, which do not share a common language. We denote with $\ell_s = (\lambda_s, \gamma_s, \mu, \mathcal{O}, \mathcal{X}, \mathcal{A})$ the source language of the encoder and with $\ell_t = (\lambda_t, \gamma_t, \mu, \mathcal{O}, \mathcal{X}, \mathcal{A})$ the target language of the decoder. We assume that the mismatch is only due to the language generator and interpreter, and that the encoder and decoder share the observation, semantic and action spaces as well as that the task that they have been trained to solve. Let $P^l = \{P_1^l, P_2^l, \dots, P_N^l\}$ be the partition defined by the language generator λ_l , $l \in \{s, t\}$. Also, let $\mathbb{1}_B(x)$ denote the characteristic function on the subset $B \subset \mathcal{X}$ defined such that:

$$\mathbb{1}_B(x) = \begin{cases} 1 & \text{if } x \in B, \\ 0 & \text{if } x \notin B. \end{cases} \quad (2)$$

Given a semantic symbol $x \in \mathcal{X}$, we refer to $[x]_l$ as the atom to which x belongs in the corresponding language, i.e. $\mathbb{1}_{[x]_l}(x) = 1$. As the source and target language differ, an observation $o \in \mathcal{O}$ may be mapped to symbol $x_s = \lambda_s(o)$ by the source language, and to symbol $x_t = \lambda_t(o)$ by the target language. Semantic mismatch may arise if x_s and x_t do not belong to the same atom of the target language partition². Hence, we introduce the following metric to quantify the semantic mismatch between source and target language:

Definition 1 (Semantic mismatch) *The semantic mismatch SM between a source language ℓ_s and a target language ℓ_t is*

$$\text{SM}(\ell_s, \ell_t) = 1 - \sum_{o \in \mathcal{O}} \mu(o) \int_{\mathcal{X}} \mathbb{1}_{[\lambda_t(o)]_t}(x) \mu_{\lambda_s}(x|o) dx \quad (3)$$

Intuitively, $\text{SM}(\ell_s, \ell_t)$ characterizes the probability of semantic misinterpretation between the two languages. Given an observation $o \in \mathcal{O}$, when the symbol $x_s \sim \mu_{\lambda_s}(x|o)$ output by the source language does not belong to the corresponding semantic atom of the target language, the characteristic function is $\mathbb{1}_{[\lambda_t(o)]_t}(x_s) = 0$ which leads to semantic misinterpretation at the receiver. Note that we have $0 \leq \text{SM}(\ell_s, \ell_t) \leq 1$. When $\text{SM}(\ell_s, \ell_t) = 1(0)$, $\mathbb{1}_{[\lambda_t(o)]_t}(x) = 0(1)$ for all observations, indicating that no (all) semantic meaning is correctly conveyed. It is also worth noting that $\text{SM}(\ell_s, \ell_t)$ is not symmetric.

Errors at the semantic level do not necessarily translate to errors at effectiveness level. The effectiveness mismatch is a metric centred around the actions and how they affect the performance on the task. To capture how each action affects the task, we define, inspired by the the Reinforcement Learning (RL) literature, the action value function $Q : \mathcal{A} \times \mathcal{O} \rightarrow \mathbb{R}$ which assigns to each observation-action pair a real number. The number $Q(a, o)$ is called the ‘‘value’’ of taking action a when the observed state is o and it is related to the optimality of the action. If we interpret solving a task as equivalent to optimizing an objective function, the value $Q(a, o)$ measures the effect of the decision on the objective function. The higher

²Important note: x_s and x_t can differ, i.e. $x_s \neq x_t$ but they still belong to the same atom; there is no semantic error in this case as semantic is associated to the atoms of a partition.

the value $Q(a, o)$, the better that choosing action a in the context of o is at optimizing the objective. So we should interpret Q as a function that allows to find the best action for the current observation and to evaluate how sub-optimal other actions are. Using the function Q , we can then define the effectiveness mismatch as

Definition 2 (Effectiveness mismatch) *Given source language ℓ_s and target language ℓ_t and denoting $\hat{a}(o) = \gamma_t(\lambda_s(o))$ as the interpreted action for observation o when the source encoder λ_s communicates with the target decoder γ_t , the effectiveness mismatch can be measured by*

$$\text{EM}(\ell_s, \ell_t) = 1 - \sum_{o \in \mathcal{O}} \mu(o) \frac{Q^+(\hat{a}(o), o)}{Q^+(a^*(o), o)}, \quad (4)$$

where $Q^+(a, o) = Q(a, o) - \min_{a'} Q(a', o)$ is the positive-action value function for the task and $a^*(o) = \arg \max_a Q(a, o)$. Given that $0 \leq Q^+(a, o) \leq Q^+(a^*(o), o)$ for all $o \in \mathcal{O}$ and $a \in \mathcal{A}_t$ and that μ is a probability measure on \mathcal{O} , we have $0 \leq \text{EM}(\ell_s, \ell_t) \leq 1$.

The effectiveness mismatch is a normalized measure that captures how the performance on the task is degraded due to the language mismatch. When $\text{EM}(\ell_s, \ell_t) = 0$, even if there is some language mismatch and subsequent semantic misinterpretations, the performance on the task is the optimal one. On the other hand, when $\text{EM}(\ell_s, \ell_t) = 1$, the language mismatch is such that the performance is the worst possible one. As with the semantic mismatch, the effectiveness mismatch is not symmetric and depends on the choice of source and target languages. We note that, assuming that the target language leads to a behaviour that follows Q , $\text{EM}(\ell_s, \ell_t) \neq 0$ implies that $\text{SM}(\ell_s, \ell_t) \neq 0$ (task performance degradation implies semantic misinterpretation), but $\text{SM}(\ell_s, \ell_t) \neq 0$ does not imply $\text{EM}(\ell_s, \ell_t) \neq 0$ (misinterpretation does not imply task performance degradation). This is possible because multiple actions might achieve the optimal performance at a given time step and thus incorrect action interpretation might still lead to optimal performance. This is the principal distinction between semantic and effectiveness mismatch that will motivate the different equalization strategies developed in this work.

III. SEMANTIC CHANNEL EQUALIZATION

In this section we will introduce the concept of semantic channel equalization through *measurable transformations*. A measurable transformation T is a function that transforms values measured by a probability distribution into values measured by another probability distribution. This way, by using μ_{λ_s} and μ_{λ_t} as measures over \mathcal{X} , a transformation $T : \mathcal{X} \rightarrow \mathcal{X}$ is a measurable transformation. We call this type of transformation a *semantic equalizer*. Semantic equalization between a source language ℓ_s and a target language ℓ_t is the process of transforming, for each observation o , a source semantic symbol $x_s = \lambda_s(o)$ into $\hat{x}_t = T \circ \lambda_s(o)$. We denote the set of measurable transformations of \mathcal{X} onto itself as \mathcal{T} . We can then define a language equalization,

Definition 3 (Semantic language equalization) Let $\ell = (\lambda, \gamma, \mu, \mathcal{O}, \mathcal{X}, \mathcal{A})$ be a language and $T \in \mathcal{T}$ be a semantic equalizer, the equalized version of ℓ is

$$T\ell = (T\lambda, \gamma, \mu, \mathcal{O}, \mathcal{X}, \mathcal{A}) \quad (5)$$

where $T\lambda = T \circ \lambda$ denotes the transformed language generator.

In this work we deal with equalization from the side of the encoder, hence the modified language generator $T\lambda$ in the definition of equalized language, but notice equalization can be performed at the decoder as well.

The optimal semantic equalizer is the one that minimizes either the semantic or the effectiveness mismatch between the transformed source language and the target language. It can be found by solving the optimization problem

$$T^* = \arg \min_{T \in \mathcal{T}} [f(T\ell_s, \ell_t)]. \quad (6)$$

Where $f(T\ell_s, \ell_t) = \text{SM}(T\ell_s, \ell_t)$ if the objective is to minimize the semantic mismatch and $f(T\ell_s, \ell_t) = \text{EM}(T\ell_s, \ell_t)$ if the objective is to minimize the effectiveness mismatch. We aim at finding the optimal semantic equalizer T^* . Solving the problem in Eq. (6) is challenging and a transformation that can correctly map the whole source semantic space into the target semantic space can be highly complex. Therefore, we propose to design a *codebook* of multiple low complexity transformations that, together with a selection policy, can effectively equalize the semantic channel.

A. Transformation Codebook

We focus now on languages with finite discrete action spaces, which makes the design of a finite and discrete transformation codebook feasible. In this case, we can assign to each action a semantic atom, and thus the semantic partition will also be discrete and finite. The transformation codebook we plan to design consists of multiple low-complexity linear transformation that map semantic atoms from the source language into semantic atoms of the target language. Let $\mathcal{A}_i = \{a_i^l | l \in J_i\}$ be the action set of language ℓ_i , where $J_i = \{1, 2, \dots, K_i\}$ denotes the action set index and let P_i^l be the atom of the semantic space corresponding to action a_i^l of language ℓ_i . The codebook of transformations can be designed by solving the following problem

$$\mathcal{C}_{\mathcal{T}} = \left\{ T_{i,j} = \arg \max_{T \in \mathcal{T}} [\rho_{P_i^s \rightarrow P_j^t}(T)] \forall i \in J_s, j \in J_t \right\} \quad (7)$$

where

$$\rho_{P_i^s \rightarrow P_j^t}(T) = \frac{\mu_{T\lambda_s}(T(P_i^s) \cap P_j^t)}{\mu_{T\lambda_s}(P_i^s)} \quad (8)$$

is the information transfer between source semantic atom P_i^s and target semantic atom P_j^t under the semantic equalizer T [16]. Here, we define $T(A) = \{T(x) | x \in A\}$. Since $\mu_{T\lambda_s}$ is the equalized source distribution on the semantic space, $\rho_{P_i^s \rightarrow P_j^t}(T)$ measures the probability that an element from P_i^s falls in P_j^t after being equalized by T . We also have that $0 \leq \rho_{P_i^s \rightarrow P_j^t}(T) \leq 1$.

B. Semantic and effectiveness channel equalization using optimal transport

This work leverages Optimal Transport (OT) to solve the problem in Eq. (7). OT is a method to map a set of samples from a source distribution into samples from a target distribution while minimizing the transportation cost (e.g. euclidean distance). This effectively means transforming samples from the source distribution in a way that they appear to be sampled from the target distribution. In particular, in [17, Sec. 3.2] a method to obtain a linear approximation of the optimal transportation map is presented. By using this method, we show that it is possible to compensate non linear language mismatch by operating a codebook of *low-complexity linear-transformations*. For more details on the optimal transport algorithm used to obtain the codebook we refer the reader to [14].

C. Policy for transformation selection

To complete the semantic equalization method, once $\mathcal{C}_{\mathcal{T}}$ is obtained, a policy π that selects a transformation from $\mathcal{C}_{\mathcal{T}}$ and applies it on a given semantic symbol needs to be designed. The selection policy can be designed to either minimize the semantic mismatch (Definition 1) or the effectiveness mismatch (Definition 2). For each observation $o \in \mathcal{O}$, the encoder has to select $T \in \mathcal{C}_{\mathcal{T}}$ to pre-equalize his message following a policy π that will minimize the desired metric. We will focus on finding a probabilistic policy π where $\pi(o, T)$ specifies the probability of choosing transformation T to equalize $\lambda_s(o)$. We define

$$\pi(o, T) = \arg \min_{\hat{\pi} \in \Pi} \mathbb{E}_{o \sim \mu(o)} [R(\hat{\pi}, o)] \quad (9)$$

where $R \in \{R_S, R_E\}$ is a risk measures that can either be the *semantic mismatch* risk or the *effectiveness mismatch* risk, both this risk functions will be detailed next. .

1) *Semantic mismatch risk*: To minimize the semantic mismatch, a selection policy should choose the transformation T that maps each semantic $x_s \in P_i^s$ into a corresponding target symbol $\hat{x}_t \in \{P_j^t | j \in \kappa(i)\}$. Here $\kappa: J_s \rightarrow J_t$ captures the correspondence between source and target semantic atoms. κ can be a one-to-one mapping, many-to-one or one-to-many depending on the source and target language generators, the action space and the downstream tasks to solve by both languages. In our scenario, both the source and the target language were designed to solve the same task with the same action space, so we can find a one-to-one correspondence between atoms, i.e. atoms are associated if they lead to the same action. We then define the semantic mismatch risk as

$$R_S(\pi, o) = 1 - \mathbb{E}_{T \sim \pi(o, T)} \left[\sum_{i \in J_s} \mu_{\lambda_s}(P_i^s | o) \sum_{j \in \kappa(i)} \rho_{P_i^s \rightarrow P_j^t}(T) \right] \quad (10)$$

where $\mu_{\lambda_s}(P_i^s | o)$ is the probability that the semantic symbol $x_s = \lambda_s(o)$ belongs to the semantic atom P_i^s .

2) *Effectiveness mismatch risk*: Not all semantic misinterpretations lead to the same degradation in performance, even more, it is possible that, after experiencing a semantic error, the effectiveness aspect of the system isn't affected at all. This is the main motivation behind reducing the effectiveness mismatch rather than the semantic mismatch. Denoting as $Q_t(a, o)$ the target language estimation of the true $Q(a, o)$ value function, the effective mismatch risk is

$$R_E(\pi, o) = 1 - \mathbb{E}_{T \sim \pi(o, T)} \left[\sum_{i \in J_s} \mu_{\lambda_s}(P_i^s | o) \sum_{j \in J_t} \rho_{P_i^s \rightarrow P_j^t}(T) \cdot Q_t(a_j, o) \right] \quad (11)$$

where $a_j = \gamma_t(P_j^t)$ is the action the target decoder outputs for any semantic message inside partition P_j^t . Notice that, since it relies on the estimation $Q_t(a, o)$, this equalization strategy is tied to the performance of the target language.

IV. NUMERICAL RESULTS

We test the proposed system on a language designed by RL methods to solve the environment shown in Fig. 1. In this scenario, the observation space \mathcal{O} is the state of a grid world with an agent and a treasure. The encoder λ maps each observation into $\mathcal{X} = \mathbb{R}^2$ in a deterministic manner and transmits it through a AWGN channel. The output power of the encoder is normalized so that each of the components of the semantic symbol $x = (x_1, x_2)$ is limited to 1, i.e. $x_i \leq 1, i \in \{1, 2\}$. This way, since it is the maximum and not the mean signal power which is fixed, the characteristics of the channel are defined by the Peak Signal to Noise Ratio (PSNR) rather than by the Signal to Noise Ratio (SNR). However, to simplify notation, from now on we will denote the PSNR as SNR. The decoder γ receives the noisy version of the transmitted symbol and chooses an action from $\mathcal{A} = \{\text{right, down, left, up}\}$ which is then executed by the agent. For the action taking, we will explore both a deterministic decoder and a stochastic one. The objective is to get the agent to the treasure in the least number of time steps possible. The episode finishes when the agent gets to the treasure or when the maximum number of steps (150) is achieved.

We assume that the encoder and decoder are given and that they were jointly trained in a de-centralized manner using RL techniques with the same reward signal with a SNR of 10 dB. The decoder is trained to learn $Q(a, x)$ using Deep Q-Learning (DQN) [18] which is suitable for discrete action spaces while the encoder is trained using Deep-Deterministic policy gradient (DDPG) [19] which is suitable for a continuous action space. For more details on the training we refer the reader to [11, Section IV].

During evaluation time, the decoder decision making is defined as

$$\gamma(x) = a \sim \text{SoftMax}(\beta Q(a, x)) \quad (12)$$

where β is the inverse of the well known temperature value for a soft max distribution where $\beta = 0$ corresponds to completely

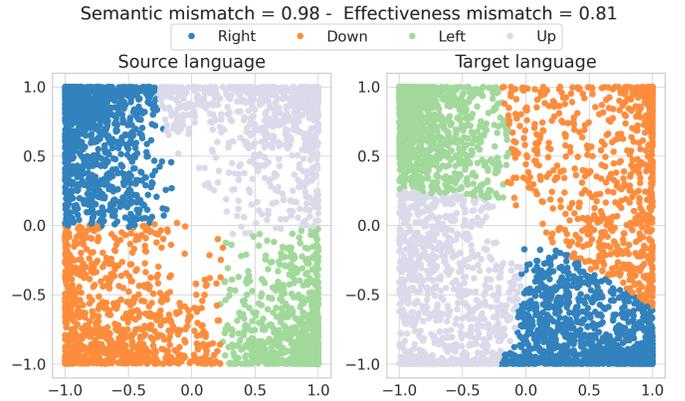


Figure 2. Mapping of the learned languages. Both encoder-decoder pairs were trained to solve the same task in identical conditions only differing in stochastic training aspects.

random decisions and the deterministic case can be recovered by setting $\beta \rightarrow \infty$. For the results we will explore both a deterministic decoder and a stochastic one with $\beta = 5$.

A. Emergent language diversity

In Fig. 2 the semantic language partition is shown for both the target and the source language. We see that each language follows a logic that divides the semantic space in different ways with a successful task competition. The training algorithm, NN architectures and training task are identical for both languages, this shows that the training results are sensible to random weight initialization and randomness in training. In particular we see that opposite actions (up-down and right-left) are mapped the furthest away allowed by the available semantic space (which is limited by the power normalization) for both languages. This characteristic suggests that, even if multiple optimal semantic space partitions are possible, there is a certain structure that allows this languages to successfully complete a task. This latent space invariability has been observed for non distributed Artificial Intelligence (AI) in multiple task [15].

B. Language equalization

The metric to evaluate the performance of the system is the episode length. The objective of the communication is for the agent to find the treasure as fast as possible, thus, lower episode length means better performance. We compare the performance of the proposed methods with other approaches both for a deterministic decoder (see Fig. 3) and a stochastic one (see Fig. 4). The evaluated communication strategies are:

No equalization: The source encoder messages are not equalized i.e. the messages are transmitted as the encoder outputs them without further processing. We see that both for a deterministic (Fig. 3) and a stochastic (Fig. 4) decoder the performance of the system is worst. It seems counter intuitive to see that for low SNR values, the performance actually improves. This is due to the fact that the errors caused by the language mismatch can sometimes be corrected by the channel noise.

Source/Target grounded communications: The encoder and decoder both adopt the source/target language, the two ideal cases with no mismatch. For the deterministic case, Fig. 3, the languages attain the optimal performance in the SNR range from 0 to 10 dB rather than SNR = ∞ . This is related to the limitation of the learned languages. Indeed, in the absence of channel noise, the encoder-decoder pairs do not always perform optimally. In fact, we have observed that the *language systematic error* causes the system to enter in loops that lead to task failing. This loops occur when the actions of the decoders bring them to a state that has already been visited before. In this situation, it is easy to see that, without any channel noise, the agent is doomed to follow the same steps as before and the system will end up transitioning between previously encountered states. Eventually, this will lead to the agents attaining the maximum time-steps and fail the task. It is the presence of noise that can break this type of infinite loops caused by the systematic error. This is further validated by the results for stochastic decoder shown in Fig. 4. When we introduce some randomness in the decision making, the performance is not degraded for increasing SNR since the infinite loops in the episode can be avoided. In this case, channel noise is actually detrimental and the episode length is minimized for infinite SNR. This suggests that, in the deterministic decision-making scenario, reducing the syntactic channel noise may result in the increase of the semantic noise, which affect the performance in achieving the goal. In contrast, when considering the stochastic decoder (Fig. 4), reducing the syntactic noise improves the goal achievement performance. These results show the strong interplay between syntactic noise and semantic noise and how they affect each other.

SM and EM oriented equalization: The proposed OT based equalization method. For SM oriented equalization, the policy that minimized the semantic mismatch risk in Eq. (10) is used. For EM oriented equalization, we employ a policy that minimizes the effectiveness mismatch risk in Eq. (11). Notice that, since it follows μ_{λ_s} , the SM oriented equalization is tied to the source language performance. On the other hand, the EM oriented equalization is tied to the target language performance by the target language Q function Q_t . In both the cases of a deterministic (Fig. 3) and a stochastic (Fig. 4) decoder, our methods are effective on minimizing the language mismatch effects on the performance since the performance is improves for all values of SNR when compared to the *no equalization* approach. However, the equalization methods perform worse than the base languages. This can be explained by the resilience of the proposed methods to channel noise, indeed, as the developed methods reduce the effect of the wireless channel noise, this enforces the systematic errors of the language and causes an increase in the average episode length. This hypothesis is further reinforced by the fact that, under error-free syntactic channel (SNR = ∞) and deterministic decoding (Fig. 3), the performance of the SM and EM oriented equalization decreases and aligns perfectly with the source and target grounded communications respectively.

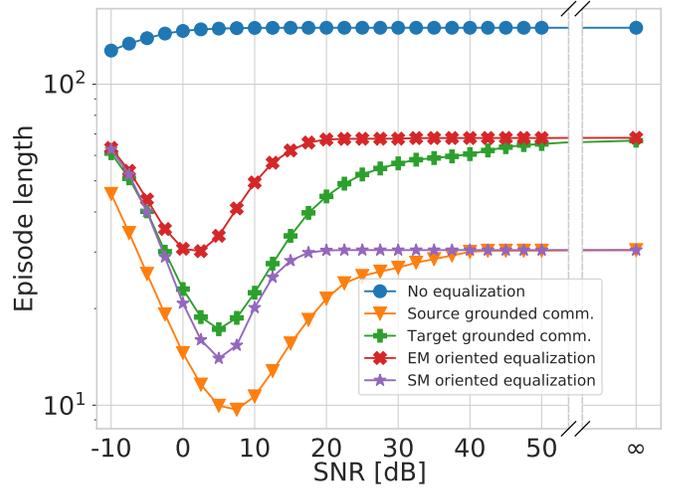


Figure 3. Average episode length (lower is better) for the different communication strategies with varying SNR for a deterministic decoder ($\beta \rightarrow \infty$ in Eq. (12)).

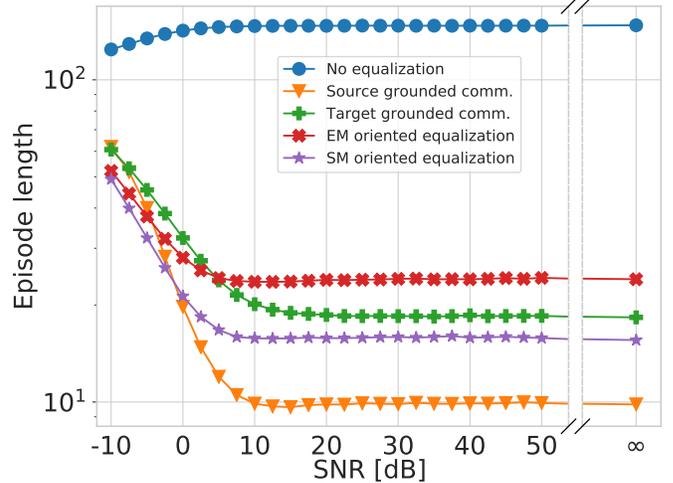


Figure 4. Average episode length (lower is better) for the different communication strategies with varying SNR for a stochastic decoder with $\beta = 5$ in Eq. (12).

C. Ablation study on β

To better characterize the effects of a stochastic decoder, we performed a study of the performance of the different methods while varying β and fixing SNR = ∞ . In Fig. 5 the results are shown. We see that increasing β has similar effects as increasing the SNR for the deterministic encoder (Fig. 3). Indeed, as β increases, the decoder becomes more and more deterministic to choose actions based on the maximum the Q -values principle (see Eq. (12)). Conversely, when β decreases, the decoder becomes more stochastic. In particular, for $\beta = 0$, all the decisions are random and thus independent from the communication, we see in this case that all the communication methods attain the same performance.

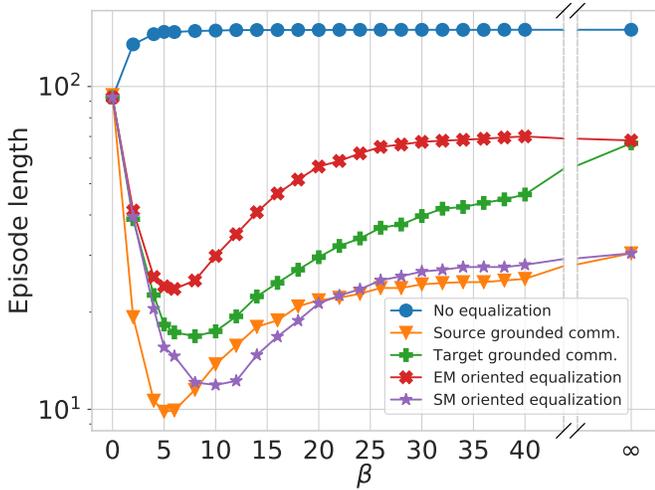


Figure 5. Average episode length (lower is better) for the different communication strategies with varying β and $\text{SNR} = \infty$.

V. CONCLUSIONS

In this paper, we analyze the consequences of semantic extraction and interpretation mismatches on goal-oriented task solving in scenarios involving heterogeneous AI agents utilizing distinct languages. At first, we establish mathematical models that address errors originating from language mismatches, encompassing both semantic and effectiveness levels, employing measurable transformations within semantic representation spaces. We show that in case of language mismatch, emergent languages might encounter systematic errors in task solving. In order to assess and address this critical issue, we present innovative metrics aimed at quantifying the influence of language mismatch across both semantic and effectiveness dimensions. Our proposed solution incorporates these metrics to devise an innovative equalization algorithm firmly grounded in optimal transport theory. This algorithm serves as a crucial component in mitigating systematic errors stemming from language mismatches. Notably, our proposed equalization methods effectively diminish language-related discrepancies, thereby enhancing the overall robustness of task execution. Furthermore, we investigate the behaviour of the system and validate the performance of our equalization methods when stochastic decision-making is used. This helps to understand the role that randomness and errors play in distributed MDP problems.

ACKNOWLEDGEMENTS

The present work was supported by the EU Horizon 2020 Marie Skłodowska-Curie ITN Greenedge (GA No. 953775),

by “6G-GOALS”, an EU-funded project, and by the French project funded by the program “PEPR Networks of the Future” of France 2030.

REFERENCES

- [1] E. Calvanese-Strinati, S. Barbarossa, J. L. Gonzalez-Jimenez, D. Ktenas, N. Cassiau, L. Maret, and C. Dehos, “6G: The next frontier: From holographic messaging to artificial intelligence using subterahertz and visible light communication,” *IEEE Veh. Technol. Mag.*, vol. 14, pp. 42–50, Sep. 2019.
- [2] Y. Xiao, G. Shi, Y. Li, W. Saad, and H. V. Poor, “Toward self-learning edge intelligence in 6g,” *IEEE Communications Magazine*, vol. 58, pp. 34–40, December 2020.
- [3] E. Calvanese-Strinati and S. Barbarossa, “6G networks: Beyond shannon towards semantic and goal-oriented communications,” *Computer Networks*, vol. 190, p. 107930, 2021.
- [4] 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, 2021.
- [5] M. Kountouris and N. Pappas, “Semantics-empowered communication for networked intelligent systems,” *IEEE Communications Magazine*, vol. 59, no. 6, pp. 96–102, 2021.
- [6] C. E. Shannon, W. Weaver, and N. Wiener, “The mathematical theory of communication,” *Physics Today*, vol. 3, no. 9, p. 31, 1950.
- [7] Y. Bar-Hillel and R. Carnap, “Semantic information,” *British Journal for the Philosophy of Science*, vol. 4, no. 14, 1953.
- [8] J. Bao *et al.*, “Towards a theory of semantic communication,” in *IEEE Network Science Workshop*, 2011.
- [9] J. Foerster, I. A. Assael, N. De Freitas, and S. Whiteson, “Learning to communicate with deep multi-agent reinforcement learning,” *Advances in neural information processing systems*, vol. 29, 2016.
- [10] A. Lazaridou, K. M. Hermann, K. Tuyls, and S. Clark, “Emergence of linguistic communication from referential games with symbolic and pixel input,” *arXiv preprint arXiv:1804.03984*, 2018.
- [11] 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, 2021.
- [12] E. Boursoulatzé, D. B. 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, 2019.
- [13] M. Sana and E. Calvanese-Strinati, “Learning semantics: An opportunity for effective 6G communications,” pp. 631–636, 2022.
- [14] M. Sana and E. Calvanese-Strinati, “Semantic Channel Equalizer: Modelling Language Mismatch in Multi-User Semantic Communications,” *arXiv preprint arXiv:2308.03789*, 2023.
- [15] L. Moschella, V. Maiorca, M. Fumero, A. Norelli, F. Locatello, and E. Rodola, “Relative representations enable zero-shot latent space communication,” *arXiv preprint arXiv:2209.15430*, 2022.
- [16] S. Sinha, U. Vaidya, and E. Yeung, “On information transfer in dynamical systems with applications in control of non-equilibrium dynamics,” in *2019 Sixth Indian Control Conference (ICC)*, pp. 326–331, 2019.
- [17] M. Perrot, N. Courty, R. Flamary, and A. Habrard, “Mapping estimation for discrete optimal transport,” in *Advances in Neural Information Processing Systems* (D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, eds.), vol. 29, Curran Associates, Inc., 2016.
- [18] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, “Human-level control through deep reinforcement learning,” *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [19] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *arXiv preprint arXiv:1509.02971*, 2015.