

Beyond Transmitting Bits: Context, Semantics and Task-Oriented Communications

Deniz Gündüz, Zhijin Qin, Inaki Estella Aguerri, Harpreet Dhillon, Zhaohui Yang, Aylin Yener, Kai Kit Wong, Chan-Byoung Chae

Abstract—Communication systems to date primarily aim at reliably communicating bit sequences. Such an approach provides efficient engineering designs that are agnostic to the meanings of the messages or to the goal that the message exchange aims to achieve. Next generation systems however can be potentially enriched by folding message semantics and goals of communication into their design. Further, these systems can be made cognizant of the context in which communication exchange takes place, providing avenues for novel design insights. This tutorial summarizes the efforts to date, starting from its early adaptations, semantic-aware and task-oriented communications, covering the foundations, algorithms and potential implementations. The focus is on approaches that utilize information theory to provide the foundations, as well as the significant role of learning in semantics and task-aware communications.

I. INTRODUCTION

Digital communication systems have been conceptualized, designed, and optimized for the main design goal of reliably transmitting bits over noisy communication channels. Shannon’s channel coding theory provides the fundamental limit on the rate of this reliable communication, with the crucial design choice that the system design remains oblivious to the underlying meaning of the messages to be conveyed or how they would be utilized at their destination. It is this disassociation from content that allows Shannon’s approach to abstract the “engineering” problem that is to replicate a digital sequence generated at one point, asymptotically error-free at another. While this approach has been tremendously successful in systems where communicating, e.g., voice and data, is the main objective, many emerging applications, from autonomous driving, to healthcare and Internet of Everything, will involve connecting machines that execute (sometimes human-like) tasks. In such applications, the goal is often not to reconstruct the underlying message exactly as in pure communications applications, but to enable the destination to make the right inference or to take the right decision and action at the right time and within the right context. Similarly, human-machine interactions will be an important component, where humans will simultaneously interact with multiple devices using text, speech, or image commands, leading to the need for similar interaction capabilities on the devices. These applications motivate the development of “semantic” and “task-oriented” communication systems. Recent advances in artificial intelligence technologies and their applications have boosted the interest and the potentials of semantics and task-oriented communications, particularly within the context of future wireless communications systems, such as 6G. Integrating semantics into system design, however, is still in its infancy with ample open problems ranging from foundational formulations to development of practical systems. This article provides an introduction to tools and advancements to date in semantic and task-oriented communications with the goal of providing a comprehensive tutorial for communication theorists and practitioners.

A. Motivation of Semantics and Task-oriented Communications

There is a growing interest in semantic and goal-oriented communication systems in the recent years. This interest is mainly driven

by new verticals that are foreseen to dominate the data traffic in future communication networks. Current communication networks are designed to serve data packets in a reliable and efficient manner without paying attention to the contents of these packets or the impact they would have at the receiver side. However, there is a growing understanding that many of the emerging applications can benefit from going beyond the current paradigm that completely separates the design of the communication networks from the source and destination of the information that flows through the network. The reason behind this trend is twofold: First, the success of many of the emerging applications, such as autonomous driving or smart city/home/factory, relies on massive datasets that enable training large models for various tasks. Hence, supporting and enabling such applications will require carrying significant amounts of traffic due to the transmission of these massive datasets and large models, which can potentially saturate the network capacity. For example, each self-driving car collects terrabytes of data each day from its many sensors, including radar, LIDAR, cameras, and ultrasonic sensors. Such data is often collected by the manufacturers to test and improve their models, generating a huge amount of traffic. Hence, the communication infrastructure cannot simply be an ignorant bit-pipe to enable intelligence at the higher layers, but must incorporate intelligence itself to make sure only the required traffic is transmitted at the necessary time and speed [1]. This will require a data-aware communication network design that has the required intelligence to understand the relevance, urgency, and meaning of the data traffic in conjunction with the underlying task. Second, unlike the current content-delivery type traffic, most of the emerging applications require extremely low end-to-end latency. For example, when delivering a video signal to a human, even in the case of live streaming, a certain level of delay is acceptable. However, when this video signal is intended to be used by an autonomous vehicle to detect and avoid potential obstacles or pedestrians on the street, even small delays may not be acceptable. On the other hand, to carry out such a task does not require the vehicle to receive the video sequence with the highest fidelity that is often considered when serving a human receiver. The vehicle is only interested in the content that is relevant for the task at hand. Again, it is essential to incorporate intelligence into the communication system design to extract and deliver the task-relevant information in the fastest and most reliable manner.

B. What is Semantics? A Historical Perspective

Semantics has been studied for centuries in the context of different disciplines, such as philosophy, linguistics, and cognitive sciences, to name a few. It is a highly complex, and to some extent, controversial topic, which is very difficult, if not impossible, to define in a concise manner that would be widely acceptable. Very briefly, semantics can be defined as the study of “meaning”, but maybe not surprisingly, it means different things in different areas. It is closely connected to *semiotics*, the study of *signs*. All communication systems are built upon signs, and a language can be broadly defined as a system of signs and rules. The rules applied to signs are divided into three

categories [2]: 1) *Syntax*: studies signs and their relations to other signs; 2) *Semantics*: studies the signs and their relation to the world; and 3) *Pragmatics*: studies the signs and their relations to users.

In this classification, syntax is only concerned with the signs and their relations, and according to Cherry, “treats language as a calculus” [3]. Semantics is built upon syntax, and its main goal is to understand the relations between signs and the objects to which they apply, *designata*. According to Chomsky, syntax is independent of semantics [4]. Through his now famous sentence “Colorless green ideas sleep furiously,” Chomsky argued that it is possible to construct grammatically consistent but semantically meaningless phrases, hence the separation between syntax and semantics. Pragmatics, on the other hand, is the most general of the three, and considers context of communication; that is, takes into account all the personal and psychological factors (in human communications) when considering the impact of a sign on *designata*.

Possibly inspired by the above classification, in his accompanying article to Shannon’s Mathematical Theory of Communication in [5], Weaver identified three levels of communication problems, in parallel with the three categories above. According to Weaver, Level A deals with the *technical problem*, and tries to answer the question “How accurately can the symbols of communication be transmitted?”. Level B instead deals with the *semantic problem*, which asks “How precisely do the transmitted symbols convey the desired meaning?”. Finally, the third level, Level C, corresponds to the effectiveness problem, and asks “How effectively does the received meaning affect conduct in the desired way?”. Similarly to syntax, the engineering problem can be considered as the syntax in communications, dealing only with their relations and how they are transmitted over a communication channel.

Similarly to Chomsky’s strict separation between syntax and semantics, Shannon’s information theory deals exclusively with the engineering problem, ignoring the meaning of the symbols transmitted. Indeed, in his seminal paper [6], Shannon explicitly states the following: “The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have *meaning*; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is one selected from a set of possible messages. The system must be designed to operate for each possible selection, not just the one which will actually be chosen since this is unknown at the time of design.”

Despite Shannon’s clear indication, many researchers at the time were excited about using Shannon’s statistical theory to explain or measure semantics. For example, in [7], Wiener wrote “The amount of meaning can be measured. It turns out that the less probable a message is, the more meaning it carries, which is entirely reasonable from the standpoint of common sense.” Weaver, on the other hand, although explicitly excluding semantics from Shannon’s theory, still argues it has implications for Level B and Level C problems. He writes [5] “[Shannon’s theory] although ostensibly applicable only to Level A problems, actually is helpful and suggestive for the level B and C problems.”. And later adds: “Thus when one moves to levels B and C, it may prove to be essential to take account of the statistical characteristics of the destination. One can imagine, as an addition to the diagram, another box labeled “Semantic Receiver” interposed between the engineering receiver (which changes signals to messages) and the destination. This semantic receiver subjects the message to a second decoding, the demand on this one being that it must match the statistical semantic characteristics of the message to the statistical semantic capacities of the totality of receivers, or of that subset of

receivers which constitute the audience one wishes to affect.”

Not everybody was in the same opinion as Wiener and Weaver regarding the potential application of Shannon’s information theory to semantics. In [8], Bar-Hillel and Carnap argue that the statistical theory of information conceived by Shannon cannot be applied to study semantics. They also express their dissatisfaction with such attempts: “Unfortunately, however, it often turned out that impatient scientists in various fields applied the terminology and the theorems of Communication Theory to fields in which the term “information” was used, presystematically, in a semantic sense, that is, one involving contents or *designata* of symbols, or even in a pragmatic sense, that is, one involving the users of these symbols.”

Based on Shannon’s explicit statement of excluding semantics from his theory, and the ensuing discussion by Carnap and Bar-Hillel, today, many authors argue that the classical Shannon theory cannot handle the semantics related aspects of communication systems. In Shannon’s channel coding theorem, the goal is to convey the maximum number of bits through a communication channel in a reliable manner, where “reliable” means that the transmitted bit sequence must be reconstructed at the receiver with an arbitrarily low probability of error. Here, each bit is assumed to be equally likely, and what the receiver does with these bits is not relevant for the channel capacity. But, not all information sources generate sequences of equally likely bits. This is obviously not the case in a text in any language. Shannon also looked at such information sources, and showed in his *source coding theorem* [6] that any information source can be compressed into equally likely messages at a rate at least at the entropy of the information source (assuming the information source generates independent symbols from an identical distribution). Even though Shannon explicitly studied the entropy of the English language as an example [9], like channel coding theory, his source coding theory does not deal with the meaning of the words. From the point of Shannon theory, an information source generating messages from the set {I love you, I miss you, I can’t stand you} is the same as the one generating messages from the set {1, 2, 3} as long as the messages come from the same probability distribution.

However, while Shannon’s theory did not deal with the meaning of these messages, it did not ignore the possibility of imperfect reconstructions. Even in his seminal work [6], which mostly focused on the reliable transmission of sources, Shannon highlighted that the exact transmission of continuous sources would require a channel of infinite capacity, but they can be delivered within a certain fidelity criterion. He laid down the basic ideas of a rate-distortion theory in [6], although the theory is developed more rigorously only in his later work in 1959 [10]. In [6], he mentioned various fidelity measures that can be considered when transmitting continuous signals, including mean squared error, frequency weighted mean squared error, and absolute error. Shannon then states: “The structure of the ear and brain determine implicitly an evaluation, or rather a number of evaluations, appropriate in the case of speech or music transmission. There is, for example, an ‘intelligibility’ criterion in which $\rho(x; y)$ is equal to the relative frequency of incorrectly interpreted words when message $x(t)$ is received as $y(t)$. Although we cannot give an explicit representation of $\rho(x; y)$ in these cases it could, in principle, be determined by sufficient experimentation. Some of its properties follow from well-known experimental results in hearing, e.g., the ear is relatively insensitive to phase and the sensitivity to amplitude and frequency is roughly logarithmic.” Here, Shannon uses the term “evaluation” to refer to different fidelity measures. He clearly makes a reference to reconstruction measures that go beyond recovering a sequence of bits, and allow a certain level of reconstruction error as long as that is within the intelligibility of the receiver, e.g., the ear or the brain. One can further argue that Shannon already hints

towards data-driven evaluation of the fidelity of a reconstruction in accordance with the machine learning approaches widely used today.

In this paper, our goal is to provide a comprehensive overview of basic Shannon theoretic concepts in relation to semantic and task-oriented communications. While we will provide a summary of semantic information theory of Bar-Hillel and Carnap and some of its future extensions as well, resolving the discussion regarding the applicability of Shannon’s statistical information theory to study semantics, particularly in the context of linguistics, is out of the scope of this paper. Indeed, our main argument is that if we are given a certain distortion measure on the pairs of transmitted and reconstructed messages, the problem then falls into the realm of information and coding theory. On the other hand, in most cases, particularly for natural languages, characterizing such a distortion measure can be extremely difficult if not impossible. Modern machine learning techniques, particularly those in the area of natural language processing (NLP), can provide such distortion measures, which are becoming increasingly accurate and useful, at least from an engineering standpoint. Therefore, our interpretation and application of semantics to communication systems is closer to Weaver’s, where semantics is simply a fidelity measure imposed by the underlying information source, captured by the concept of “semantic noise” [5].

C. Relevant Surveys and Our Contributions

Given the increasing interest in semantic and task-oriented communications, it is not surprising to note that there are quite a few existing surveys and tutorial articles focusing on this general topic already. We discuss them briefly and then list our main contributions. The authors in [11] introduced three kinds of semantic communications, human-to-human, human-to-machine, and machine-to-machine communications. The key performance indicators and system design for semantic learning mechanisms over future wireless communications were further pointed out in [12]. The goal-oriented signal processing was investigated in [13], which included the graph-based semantic language and representation of the semantic information. Besides, the architectures of semantic communications for artificial intelligence-assisted wireless networks were investigated in [14]–[16]. In [17], the technical contents and application scenarios were discussed for the intelligent and efficient semantic communication network design. The recent work in [18] discussed principles and challenges of semantic communications enabled by deep learning. Compared with the above works [11]–[18], the main focus of this paper is to provide a comprehensive introduction to semantic- and task-oriented communications through an information-theoretic viewpoint. In other words, it will be our intention to ground everything discussed in this paper in relevant information-theoretic principles.

In addition to providing a comprehensive survey of semantic and task-oriented communication systems, the main ingredients of our paper are listed below.

- We point out the differences among existing definitions of semantic entropy. We further introduce how knowledge graph based semantics can be applied to and benefit a wide variety of common tasks.
- We provide the basic information theoretic concepts for semantic and task-oriented communications. For instance, we show how the rate-distortion theory can capture the semantic distortion measure and explore the connection between information bottleneck and goal-oriented compression. To convey the class information to the receiver, rate-limited remote inference theory is discussed.
- Machine learning techniques for semantic- and task-oriented communications are discussed in two phases separately, i.e.,

training phase and prediction phase. The solutions to information bottleneck problem and remote training problem are presented.

- We introduce the information theoretic concepts for semantic and task-oriented transmission over noisy channels from the viewpoint of joint source and coding (JSCC) approach. Some practical designs for goal-oriented semantic over noisy channels are further provided for text, speech, and image.
- We discuss the idea of “timing as semantics” where the relevance or value of information is in time. Connections of this idea with the related concept of age of information (AoI) are explored by discussing a general real-time remote tracking or reconstruction problem from the rate-distortion viewpoint such that the semantic information is contained in the timing of the source samples.
- We present an effective communication framework, corresponding to the Level C communication problem put forth by Weaver. This results in a context-dependent communication paradigm, where the same message may have a different affect on the receiver depending on the context.

The paper is organized as follows. We first provide an overview of information-theoretic foundations for semantic and task-oriented communications in Section II. Then, in Section III, relevant machine learning techniques for semantic and task-oriented communications are introduced in detail. The idea of semantic information theory is further discussed in Section IV and the joint source-channel coding approach is provided in Section V. As an instance of task-oriented communications, an important class of problems where the metric relates the freshness of information along with connections to semantics is discussed in Section VI. Finally, Section VII points out the conclusions and future directions.

II. A SEMANTIC INFORMATION THEORY

To extend the engineering approach Shannon proposed to include semantics, a number of researchers have started to work on a theory of semantic communications soon after Shannon’s work. The concept of semantic entropy has been foreseen likely to play a significant role in developing a framework that considers semantics. This notion is used to quantify semantic information of the source. Until now, a number of definitions of semantic entropy have been proposed from different perspectives. A commonly-agreed upon notion is yet to emerge however.

A. Semantic Entropy

In 1952, Carnap and Bar-Hillel [19] first explicated the concept of semantic entropy of a sentence within a given language system, and provided a way to measure it as follows:

$$H(s, e) = -\log c(s, e), \quad (1)$$

where $c(s, e)$ is the degree of confirmation of sentence s on the evidence e , given by

$$c(s, e) = \frac{m(e, s)}{m(e)}. \quad (2)$$

Here, $m(e, s)$ and $m(e)$ represent the logical probability of s on e and that of e , respectively.

Different from the logic-based definition, Venhuizen *et al.* [20] derived the semantic entropy based on a language comprehension model in terms of the structure of the world instead of the probabilistic structure of the language, which can be expressed as

$$H(\mathbf{v}_t) = -\sum_{\mathbf{v}_M \in \mathcal{V}_M} P(\mathbf{v}_M | \mathbf{v}_t) \log P(\mathbf{v}_M | \mathbf{v}_t) \quad (3)$$

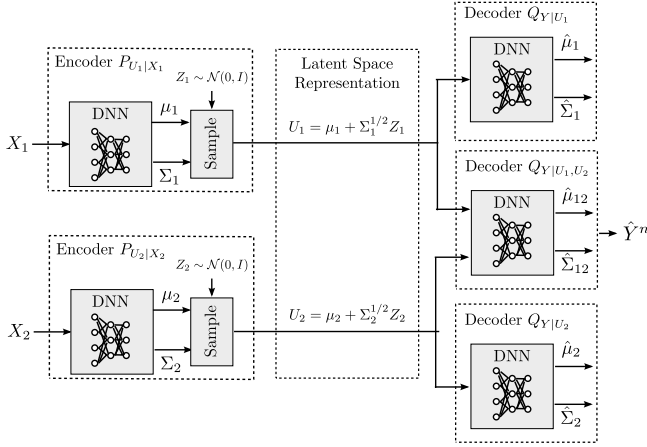


Fig. 1. Example parameterization of the distributed variational information bottleneck method using neural networks for $K = 2$ for encoders parameterized by Gaussian distributions, $P_{\theta}(u_k|s_k) = N(s_k; \mu_k^{\theta}(s), \Sigma_k^{\theta}(s))$.

where $\mathcal{V}_{\mathcal{M}} = \{\mathbf{v}_M | \mathbf{v}_M(i) = 1 \text{ iff } M_i = M \text{ and } M \text{ is a unique model in } \mathcal{M}\}$ is the set of meaning vectors that identify unique models in \mathcal{M} , \mathcal{M} is the set of models and reflects the probabilistic structure of the world, and $P(\mathbf{v}_M | \mathbf{v}_t)$ is the conditional probability of \mathbf{v}_M given \mathbf{v}_t . This comprehension-centric notion of semantic entropy depends on both linguistic experience and world knowledge and quantifies the uncertainty with respect to the whole meaning space.

Apart from the language system, the semantic entropy for intelligent tasks has also been studied. Melamed [21] proposed an information-theoretic method for measuring the semantic entropy in translation tasks by using translational distributions of words in parallel text corpora. The semantic entropy of each word w is given by

$$\begin{aligned} H(w) &= H(T|w) + N(w) \\ &= - \sum_{t \in T} P(t|w) \log P(t|w) + P(NULL|w) \log F(w) \end{aligned} \quad (4)$$

where $H(T|w)$ is the translational inconsistency of a source word w , T represents the set of target words, $N(w)$ denotes the contribution of null links of w , and $F(w)$ is the frequency of w . Additionally, for classification tasks, Liu *et al.* [22] defined the semantic entropy by introducing the membership degree in axiomatic fuzzy set theory. Let $\mu_{\varsigma}(x)$ denote the membership degree of the data sample x . The authors first obtained the matching degree, which characterizes the semantic entropy for data samples in class $C_j, j \in \{1, 2, \dots, m\}$ on semantic concept ς , as

$$D_j(\varsigma) = \frac{\sum_{x \in \mathcal{X}_{C_j}} \mu_{\varsigma}(x)}{\sum_{x \in \mathcal{X}} \mu_{\varsigma}(x)}, \quad (5)$$

where \mathcal{X}_{C_j} is the set of data for class C_j , and \mathcal{X} is the data set of all classes. According to the matching degree, the semantic entropy of class C_j on ς is defined as

$$H_{C_j}(\varsigma) = -D_j(\varsigma) \log_2 D_j(\varsigma). \quad (6)$$

Further, the semantic entropy of concept ς on \mathcal{X} can be obtained by

$$H(\varsigma) = \sum_{j=1}^m H_{C_j}(\varsigma). \quad (7)$$

Based on this definition, the optimal semantic descriptions of each class can be obtained and the uncertainty in designing the classifier is minimized.

In contrast to the aforementioned definitions that are specific to a single task, Chattopadhyay *et al.* [23] are working on exploring an information-theoretic framework to quantify semantic information for any tasks and any types of source. They defined the semantic entropy as the minimum number of semantic queries about the data X whose answers are sufficient to predict the task Y , which can be expressed as

$$\begin{aligned} H_Q(X; Y) &\triangleq \min_E \mathbb{E}_X \left[\left\| \text{Code}_Q^E(X) \right\| \right] \\ \text{s.t. } &P(y | \text{Code}_Q^E(x)) = P(y|x), \quad \forall x, y \end{aligned} \quad (8)$$

where $\text{Code}_Q^E(x)$ denotes the query vector extracted from X with the semantic encoder E . From (8), in order to obtain the semantic entropy, one needs to find the optimal semantic encoder, which has the ability to encode X into the minimal representation that can accurately predict the task. Currently, many methods have been utilized for measuring the semantic entropy, such as semantic information pursuit and variational inference. This direction is still in early stages and need to be further investigated.

In summary, there are significant differences among existing definitions of semantic entropy as each of them is based on the properties of its own concerned task. Although the last definition can be extended to different tasks, finding the optimal semantic encoder is as challenging as obtaining the semantic entropy. Hence, a unifying definition (as in the case of Shannon entropy) does not exist for semantic entropy.

B. Knowledge Graph for Semantic Communications

Fig. 2 presents an example of a knowledge graph, which represents a network of real-world entities, i.e., objects, events, situations, or concepts, and illustrates the relationship between them. This information is usually stored in a graph database and visualized as a graph structure. A knowledge graph is made up of three main components: nodes, edges, and labels. Any object, place, or person can be a node and an edge defines the relationship between the nodes. Knowledge graph embedding is to embed components of a knowledge graph including entities and relations into continuous vector spaces, so as to simplify the manipulation while preserving the inherent structure of the knowledge graph [24]. It can benefit a variety of downstream tasks and hence has quickly gained massive attention. In this following, we make the review of the semantic matching models for knowledge graph. Then, we introduce how knowledge graph based semantics can be applied to and benefit a wide variety of downstream tasks, such as data integration, recommendation systems, and so forth. Subsequently, we present the analysis and framework of knowledge graph based semantic and its applications in the semantic communication systems.

1) *Semantic Matching Model for Knowledge Graph*: The knowledge graph techniques can be roughly categorized into two groups: translational distance models and semantic matching models. The former use distance-based scoring functions, and the latter similarity-based ones. In the following, we introduce the semantic matching models. In particular, they exploit similarity-based scoring functions and measure plausibility of facts by matching latent semantics of entities and relations embodied in their vector space representations. The authors in [25] associate each entity with a vector to capture its latent semantics. Each relation is represented as a matrix which models pairwise interactions between latent factors. The authors in [26] propose the semantic matching energy method, which conducts

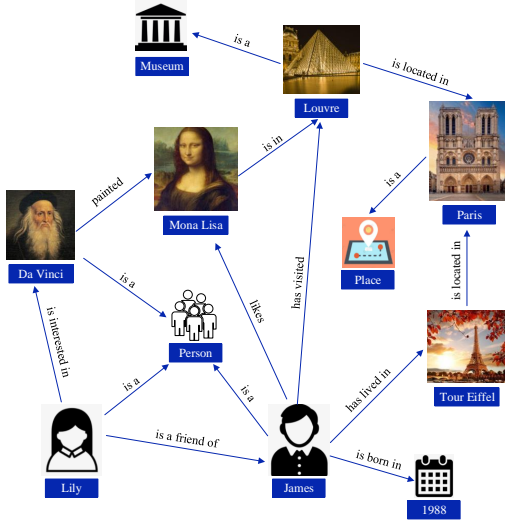


Fig. 2. The illustration of the knowledge graph.

semantic matching using neural network architectures. It first projects entities and relations to their vector embeddings in the input layer. The relation is then combined with the entity to get the score of a fact. Furthermore, the neural association model has been developed in [27] to conduct semantic matching with a deep architecture.

2) *Applications of Knowledge Graph Based Semantic*: The knowledge graph embedding is a key technology for solving the problems in knowledge graph. The authors in [28] propose a novel knowledge graph embedding method, which translates and transmits multi-directional semantics: (i) the semantics of head/tail entities and relations to tail/head entities with nonlinear functions and (ii) the semantics from entities to relations with linear bias vectors. The knowledge graph based semantic has also been employed in data integration [29], recommendation systems [30], and real-time ranking [31]–[33]. In [29], the authors devise a semantic data integration approach that exploits keyword and structured search capabilities of Web data sources. Resulting knowledge graphs model semantics or meaning of merged data in terms of entities that satisfy keyword queries, and relationships among those entities. It semantically describes the collected entities, and relies on semantic similarity measures to decide on relatedness among entities that should be merged. The authors in [30] incorporate both word-oriented and entity-oriented knowledge graphs to enhance the data representations, and adopt mutual information maximization to align the word-level and entity-level semantic spaces. In [31], a novel kind of knowledge representation and mining system has been proposed, which is referred to the semantic knowledge graph. It provides a layer of indirection between each pair of nodes and their corresponding edge, enabling edges to materialize dynamically from underlying corpus statistics. In [32], an entity-duet neural ranking model has been proposed, which introduces knowledge graphs to neural search systems and represents queries and documents by their words and entity annotations. The semantics from knowledge graphs are integrated in the distributed representations of their entities, while the ranking is conducted by interaction-based neural ranking networks. Furthermore, the translation-based embedding models attempt to translate semantics from head entities to tail entities with the relations and infer richer facts outside the knowledge graph. In [33], a novel ranking technique that leverages knowledge graph has been proposed. Analysis of the query log from the academic search engine reveals that a major error source is its inability to understand the meaning of research concepts in queries.

To address this challenge, the authors propose to represent queries and documents in the entity space and ranks them based on their semantic connections from their knowledge graph.

3) *Analysis and Framework of Knowledge Graph Based Semantic*: In [34], the authors introduce the semantic property graph for scalable knowledge graph analytics. To enhance the input data, the authors in [35] propose the framework of relevant knowledge graphs for machine learning recommendation and community detection, which improves both accuracy and explainability. In [36], the authors propose an iterative framework that is based on probabilistic reasoning and semantic embedding. The authors in [37] propose a semantic representation method for knowledge graph, which imposes a two-level hierarchical generative process that globally extracts many aspects and then locally assigns a specific category in each aspect for every triple. Because both the aspects and categories are semantics-relevant, the collection of categories in each aspect is treated as the semantic representation of this triple.

4) *Semantic Communication Systems Driven by Knowledge Graph*: In [38], a cognitive semantic communication framework is proposed by exploiting knowledge graph. Moreover, a simple and interpretable solution for semantic information detection is developed by exploiting triples as semantic symbols. It also allows the receiver to correct errors occurring at the symbolic level. Furthermore, the pre-trained model is fine-tuned to recover semantic information, which overcomes the drawback that a fixed bit length coding is used to encode sentences of different lengths.

III. INFORMATION THEORETIC FOUNDATIONS OF SEMANTIC- AND TASK-ORIENTED COMMUNICATIONS

As we have mentioned, despite the many efforts to define a semantic information measure, none of the aforementioned attempts have resulted in a widely accepted definition, or had an impact on operational performance similarly to Shannon's information theory had on communication systems. Therefore, we will take a more statistical approach to semantics, and either treat it as a given distortion measure, or consider data-driven approaches to acquire it. In the current section, we will provide some of the basic concepts in Shannon's statistical information theory, and how they can be used to study semantics in emerging communication systems. In particular, we will first review rate-distortion theory, and its characterization in various settings and under different distortion measures, and how they can capture semantic or goal-oriented communications. In this section, we will focus on rate-limited error-free communications; that is, we will mainly treat the semantic/ task-oriented compression problem. Semantic- and task-oriented transmission over noisy channels will be considered in Section V.

A. Rate-Distortion Theory

In [10], Shannon expands the above ideas put forth in his seminal work by formally defining the problem of lossy source transmission, which laid down the principles of *rate-distortion theory*. Note that, when compressing a single source sample we have a *quantization* problem, which is equivalent to analog-to-digital conversion when the source samples come from a continuous alphabet. This can also be treated as a clustering problem with a specified fidelity criterion. Shannon showed that, similarly to the channel coding theorem, it is more efficient, in terms of bits per source sample, if many independent samples of the same source distribution can be compressed together.

Let $s^m \in \mathcal{S}^m$ denote m independent source samples distributed according to $p_S(s)$, i.e., $p(s^m) = \prod_{i=1}^m p_S(s_i)$. Let $\hat{s}^m \in \hat{\mathcal{S}}^m$ denote

the reconstruction of the receiver, where $\hat{\mathcal{S}}$ is the reconstruction alphabet, not necessarily same as the source alphabet. The goal is to minimize the distortion between s^m and \hat{s}^m under some given distortion (fidelity) measure, $d : \mathcal{S}^m \times \hat{\mathcal{S}}^m \rightarrow [0, \infty)$, which assigns a certain distortion, or equivalently, a quality metric, for every pair of source and reconstruction sequences. Shannon considered single-letter additive distortion measures, that is, $d(s^m, \hat{s}^m) = 1/m \sum_{i=1}^m d(S_i, \hat{S}_i)$ for a single-letter distortion measure $d : \mathcal{S} \times \hat{\mathcal{S}} \rightarrow [0, \infty)$.

The goal in lossy source coding is to represent the source sequence with as few bits as possible, measured in *bits per source sample* (*bpss*), while guaranteeing a certain average distortion target. A $(m, 2^{mR})$ lossy source code consists of an *encoder-decoder* pair, where the encoder $f^{(m)} : \mathcal{S}^m \rightarrow [1 : 2^{mR})$ maps each m -length source sequence $s^m \in \mathcal{S}^m$ to an index $w(s^m) \in [1 : 2^{mR})$, and the decoder $g^{(m)} : [1 : 2^{mR}) \rightarrow \hat{\mathcal{S}}^m$ maps each index $w \in [1 : 2^{mR})$, to an estimated reconstruction sequence $\hat{s}^m(w) \in \hat{\mathcal{S}}^m$. The collection of all reconstruction sequences $\{\hat{s}^m(1), \dots, \hat{s}^m(2^{mR})\}$ forms the codebook, which is assumed to be shared between the encoder and decoder in advance.

The average distortion of a $(m, 2^{mR})$ code is given by

$$\mathbb{E}[d(S^m, \hat{S}^m)] = \sum_{s^m \in \mathcal{S}^m} p(s^m) d(s^m, \hat{s}^m(w(s^m))), \quad (9)$$

where the expectation is taken over the source distribution.

For a given source distribution $p_S(s)$ and distortion measure $d(s, \hat{s})$, we say that a rate-distortion pair (R, D) is *achievable* if there exist a sequence of $(m, 2^{mR})$ codes that satisfy

$$\limsup_{m \rightarrow \infty} \mathbb{E}[d(S^m, \hat{S}^m)] \leq D. \quad (10)$$

The rate-distortion function $R(D)$ of source S under the single-letter distortion measure $d(\cdot, \cdot)$ is defined as the infimum of rates R for which (R, D) is achievable.

For such single-letter additive distortion measures, Shannon provided a single-letter characterization of the optimal rate-distortion function.

Theorem 1. (*Shannon's Lossy Source Coding Theorem, [10]*) *The rate-distortion function for source S and distortion measure $d(\cdot, \cdot)$ is given by*

$$R(D) = \min_{p(\hat{s}|s): \mathbb{E}[d(S, \hat{S})] \leq D} I(S; \hat{S}). \quad (11)$$

We can argue that Shannon's rate distortion function quantifies the communication rate required to convey sample-level semantic information when many source samples can be compressed jointly. Here, our assumption is that the semantic relevance of reconstructing source sample s as \hat{s} at the receiver is quantified by the prescribed distortion measure $d(s, \hat{s})$. In the context of text, this may refer to reconstructing a certain sentence in a manner that preserves its core meaning. In the context of music or image, it may mean, as argued by Shannon, preserving the intelligibility of the source signal, for example by conveying only the most audible or most distinguishable frequencies, which is the approach followed by modern audio and image compression standards.

We note, however, that in most cases we apply compression algorithms on a single sample, e.g., a single image, a single video sequence, or a single text file. In these cases, samples correspond to pixels or patches in an image, frames in a video, and letters or words in the text. Such samples are often highly correlated, and additive distortion measures may not preserve the semantics of the overall content. For example, preserving word level similarity, e.g., replacing words with their more common synonyms in its reconstruction, or generating an image with a low pixel-level mean squared error (MSE) may not lead to a good quality semantic compression. For text,

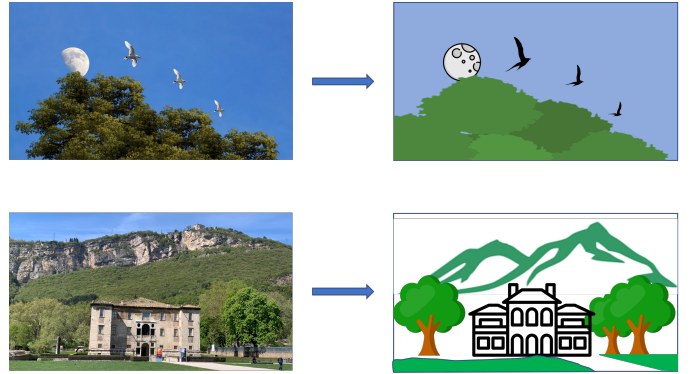


Fig. 3. Examples of semantic image compression.

semantic compression requires a much deeper understanding of the semantics of the underlying language. For images, alternative quality measures have been proposed that would provide a better image level semantic reconstruction. For example, structural similarity index measure (SSIM) or multi-scale SSIM (MS-SSIM) [39] have been introduced to measure the perceived quality of images and videos by incorporating luminance masking and contrast masking terms into the distortion measure, providing perceptually more satisfactory reconstructions. More recently, a *rate-distortion-perception trade-off* has been introduced in [40], in which the perceptual quality of an output image is defined as ‘the extent to which it is perceived as a valid (natural) sample’. We note that, while the distortion of image \hat{s}^m is defined with respect to the source image s^m , perceptual quality is defined as a property of the reconstruction. The perceptual quality of an image is often defined as the Kullback-Leibler divergence between the distribution of the sample and the statistics of natural images [41]. More recently image reconstruction techniques employing generative adversarial networks (GANs) [42] also employ similar distance measures between the reconstructed image and the statistics of the images in the training dataset [43], [44], such as the Jensen-Shannon divergence, the Wasserstein distance [45], or an f-divergence [46]. GANs train a *discriminator* that tries to distinguish the reconstructed image from the images in the dataset, forcing the decoder to generate realistic looking images. By taking the perceptual quality into account when reconstructing a compressed image, we aim not only to reproduce the original image with the highest fidelity, but also to reconstruct a more natural realistic image, which preserves the semantics of the underlying distribution.

Despite trying to preserve the overall semantics of the image, the methods above do not use the image semantic information in the compression stage. On the other hand, semantic information can be used to provide a more efficient compression algorithm, or to achieve a better quality reconstruction [44], [47], [48]. See, for example, the images in Fig. 3. By only extracting and transmitting semantic information, e.g., the objects in the original image and their general layout, the output image can be reconstructed by simply including a generic representative of each of the objects in the image. Hence, it is possible to convey the image at a level to enable semantic reasoning about the image, albeit not reliably at the pixel level. A selective generative compression method is proposed in [44], which generates parts of the image from a semantic label map, which can be obtained using a semantic segmentation network [49], [50]. These parts of the image are fully synthesized rather than being reconstructed based on the original image. The rest of the image is generated using a conditional generative adversarial network (cGAN) [51]. A deep semantic segmentation-based layered image compression scheme is

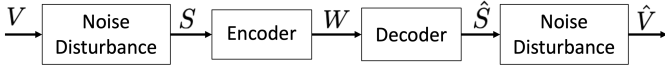


Fig. 4. Remote source-compression problem where the encoder observes a noisy version S of the remote source V through a random disturbance $p_{S|V}$, and the reconstruction \hat{S} goes through another random transformation $p_{\hat{V}|\hat{S}}$. The goal is to guarantee a certain fidelity requirement between V and \hat{V} .

proposed in [47], where the semantic map of the input image is used to synthesize the image, while a compact representation and a residual are further encoded as enhancement layers. It is shown that this semantic-based compression approach outperforms BPG and other standard codecs in both PSNR and MS-SSIM metrics. We also add that, including the segmentation map in the bit-stream can further facilitate other downstream tasks such as image search or compression and manipulation of individual image segments.

Another semantic-based image processing approach uses scene graphs to extract not only objects within the image, but also their relationships [52]. A scene graph is a directed graph data structure consisting of the objects and their attributes as vertices, and the relations between the objects as edges. Scene graph generation typically follows three steps: i) Feature extraction, which is responsible for identifying the objects in the image; ii) Contextualization, which established contextual information between entities, and finally iii) Graph construction and inference, which generates the final scene graph using the contextual information, and carry out desired inference tasks on the graph [53]. Scene graphs are powerful tools that can encode images [52] or videos [54] using abstract semantic elements.

In the context of video coding, semantic-based compression has long been considered for very low bit-rate video compression [55]. These include motion-compensated compression, where motion vectors of pixels between two consecutive images or optical flow vectors [56] are encoded and transmitted. Alternative *object-based compression* methods have also been considered in the literature for a very long time [57], [58]. In object-based compression, each moving object in a video signal is separated from the stationary background and are conveyed to the decoder by describing their shape, motion, and content using an object-dependent parameter coding. Using this coded parameter set, each image is then synthesized at the decoder by model-based image synthesis. Although such an object-based compression approach was standardised as part of MPEG4 in the late 90s, it has not been widely adopted due to the lack of fast and reliable object and motion segmentation techniques. This approach is regaining interest in recent years due to the rapid advances in deep learning based segmentation techniques [59], [60].

In general, quantifying the semantic distortion measure for a particular information source is a formidable task. There have been many studies to understand and model semantics particularly in the context of text and natural language processing. We will go over these efforts in Section V in the context of semantic/goal-oriented transmission.

B. Goal Oriented Compression

In the conventional rate-distortion framework overviewed above, the goal is to reconstruct the source sequence s^m within a desired fidelity constraint. This framework applies to most delivery scenarios, such as the transmission of an image, video or audio source over a rate-limited channel, where the goal is to recover the original content with the highest fidelity. However, in many emerging applications, particularly those involving machine-type communications, the receiver may not be interested in the source sequence, but only in

a certain feature of it. For example, rather than reconstructing an image, the receiver may be interested in certain statistical aspects of the image, or the presence or absence of certain objects or persons in the image. This can model a goal-oriented compression scenario, where reconstructing the desired feature can represent the goal.

The desired feature in this context can be modelled as a correlated random variable $V \in \mathcal{V}$ that follows a joint distribution $p(v, s)$ with the source sample S . In this scenario, we can instead impose a distortion constraint on the feature and its reconstruction at the decoder, for a prescribed distortion measure $d(v, \hat{v}) < \infty$ for $(v, \hat{v}) \in \mathcal{V} \times \hat{\mathcal{V}}$. This problem can be interpreted as a *remote source coding problem*, which was originally studied by Dobrushin and Tsybakov in [61]. In addition to the noisy observation of the features at the encoder, they also considered another random transformation at the output of the decoder, as illustrated in Fig. 4. They showed that this generalization of the Shannon's rate-distortion problem can be reduced to Shannon's original formulation by appropriately transforming the distortion measure. Consider any pair of encoding function f and decoding function g . The average end-to-end distortion achieved by this encoder-decoder pair is given by

$$D(f, g) = \sum_{v, s, \hat{s}, \hat{v}} p_V(v) p_{S|V}(s|v) p_{\hat{S}|S}^{(f, g)}(\hat{s}|s) p_{\hat{V}|\hat{S}}(\hat{v}|\hat{s}) d(v, \hat{v}). \quad (12)$$

From the perspective of the encoder, it observes a source S with marginal distribution $p_S(s) = \sum_{v \in \mathcal{V}} p_V(v) p_{S|V}(s|v)$. Let us now consider the modified distortion measure

$$\hat{d}(s, \hat{s}) = \frac{1}{p_S(s)} \sum_{v, \hat{v}} p_V(v) p_{S|V}(s|v) p_{\hat{V}|\hat{S}}(\hat{v}|\hat{s}) d(v, \hat{v}). \quad (13)$$

Using this new distortion measure, we can rewrite the end-to-end distortion as

$$D(f, g) = \sum_{s, \hat{s}} p_S(s) p_{\hat{S}|S}^{(f, g)}(\hat{s}|s) \hat{d}(s, \hat{s}). \quad (14)$$

Therefore, the problem of minimizing the average end-to-end distortion for the remote source coding problem can be reduced to the classical source coding problem for a source with marginal distribution p_S under the modified distortion measure $\hat{d}(\cdot, \cdot)$.

The following equivalence can be generalized to the standard block coding setting. Assume now that we observe a sequence s^m and want to reconstruct the corresponding feature vector v^m , where (v_i, s_i) are i.i.d. samples from the joint distribution $p_{V, S}$. Then, the corresponding *remote rate-distortion function* can be characterized in a single letter form as given in the next theorem.

Theorem 2. (*Remote Rate-Distortion Function, [61]*) *The remote rate-distortion function for source V based on its observation S following joint distribution $p_{V, S}(v, s)$ and distortion measure $d(\cdot, \cdot)$ is given by*

$$R^{\text{remote}}(D) = \min_{p(\hat{s}|s): \mathbb{E}[\hat{d}(S, \hat{S})] \leq D} I(S; \hat{S}). \quad (15)$$

Recently, the remote rate-distortion interpretation of semantic compression is considered in [62], where the decoder wants to reconstruct both the feature vector v^m and the source vector s^m , under different distortion measures. With the above reduction, one can see that this problem trivially reduces to the Shannon rate-distortion problem with two distortion measures. Further characteristics of this rate-distortion function for a Hamming distortion measure is studied in [63]. In [64], it is shown that the optimal transmission scheme for the general model in Fig. 4 under the squared error distortion measure can be divided into two steps: the encoder first estimates the feature variable

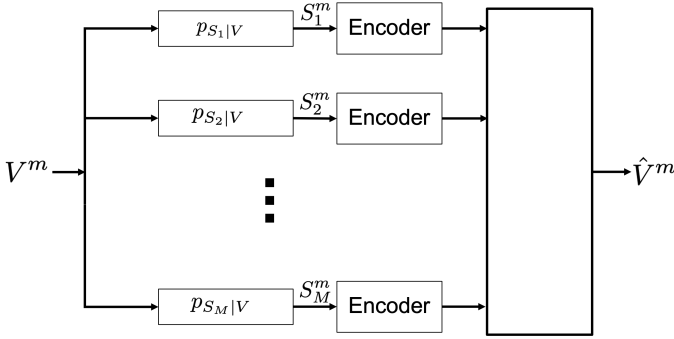


Fig. 5. The CEO problem.

v conditioned on s , and then conveys the estimated value to the decoder.

A natural extension of the above remote rate-distortion function involves multiple terminals, each observing a different noisy version of the underlying latent source V (see Fig. III-B for an illustration). This is known as the ‘‘CEO problem’’ in the literature following [65]. In this setting, a chief executive officer (CEO) is interested in estimating an underlying source sequence, V^n . She sends M agents to observe independently corrupted versions of the source sequence, where the observations S_i^n of the i -th agent are generated through the conditional distribution $p_{S_i|V}$ in an i.i.d. manner. The agents cannot communicate among each other, and each one only has a rate-constrained channel to the CEO. For a given sum rate constraint, what is the minimum average distortion the CEO can estimate S^n under a given fidelity measure $d(\cdot, \cdot) < \infty$? The special case of Gaussian source and noise statistics with squared error distortion is called the quadratic Gaussian CEO problem [66]. The problem remains open in the general case, while the rate region was characterized for the Gaussian case in [67], for logarithmic loss distortion for discrete sources in [68] and vector Gaussian sources in [69].

C. Context as Side Information

When the communication of an underlying source signal is considered, additional information that is correlated with this desired source variable available to the transmitter and the receiver can be leveraged to reduce the rate requirements. Consider for example a surveillance camera in a house recording a video and forwarding the recording to a remote node which aims at detecting the presence of intruders. Depending on the hour of the day, the illumination of the image will be different. The context in which the information is being obtained, e.g., the time, or weather conditions, could be exploited to improve the video compression quality. Images of the same scene recorded by other cameras can also serve as context information.

From an information theoretical perspective, contextual information can be modelled as side information. The problem of lossy source coding when common correlated side information is available both at the encoder and the decoder was studied by Gray in [70]. By characterizing the rate-distortion function for this problem, it was shown that the rate required to achieve a prescribed distortion could be reduced by exploiting the side information available. Interestingly, the rate required to achieve a particular distortion is also reduced when the correlated side information is only available at the decoder, and the encoder has to compress without knowledge of the realization of the side information available, as long as its distribution is known. This problem was studied in [71], and the corresponding single-letter rate-distortion function is referred to as the Wyner-Ziv rate-distortion function.

Let T^m denote the side information sequence observed by the receiver that is correlated with the source S^m . In particular, let (S_i, T_i) be i.i.d. samples jointly distributed with $p(S, T)$. The Wyner-Ziv rate distortion function for source coding problem with side information at the decoder is characterized as given in the following theorem.

Theorem 3. (Lossy Compression with Side Information at the Decoder, [71]) *The rate-distortion function for source S with side information T available only at the decoder following joint distribution $p_{S,T}(s, t)$ and distortion measure $d(\cdot, \cdot)$ is given by*

$$R_{S|T}^{\text{WZ}}(D) = \min I(S; U|T), \quad (16)$$

where the minimum is over conditional distributions $p_{U|S}$ with $|\mathcal{U}| \leq |\mathcal{S}| + 1$ and functions $g : \mathcal{U} \times \mathcal{Y} \rightarrow \hat{\mathcal{S}}$ such that $\mathbb{E}[d(S, \hat{S})] \leq D$.

In Wyner-Ziv coding, typical source codewords are split into bins, and only the bin index is forwarded to the decoder. By leveraging the side information, the receiver is able to identify the corresponding compression codeword within the selected bin. Wyner-Ziv coding has been exploited in image and video compression [72], [73]. While in general having the side information available at both the encoder and the decoder is beneficial, for some source-distortion measure pairs, e.g., Gaussian sources under squared-error distortion, it is known that having the side information available only at the decoder does not result in any performance loss [71]. Side information can also be considered for the remote source coding problem since any remote compression problem is equivalent to a standard source coding problem with a new distortion measure as in (14). Side information can also be incorporated into multi-terminal source coding problems, e.g., [74]. For example, in the CEO setting in Fig. III-B, letting one of the observers to have a link to the CEO with sufficiently large rate allows conveying this observation perfectly, which then acts as side information.

D. The Information Bottleneck and Goal Oriented Compression

The information bottleneck (IB) was introduced by Tishby et al. [75] as a methodology for extracting the information that a variable $S \in \mathcal{S}$ provides about another one $V \in \mathcal{V}$ that is of interest and not directly observable by mapping S into a representation $U \in \mathcal{U}$, as shown in Figure 4. Specifically, the IB method consists in finding the mapping $P_{U|S} : \mathcal{S} \rightarrow \mathcal{U}$ that, given S , outputs the representation U that is maximally informative about V , such that the mutual information $I(U; V)$ is maximized, while being minimally informative about S , i.e., such that the mutual information $I(U; S)$ is minimized. Here, $I(U; V)$ is referred to as the *relevance* of U and $I(U; S)$ is referred to as the *complexity* of U , where complexity is measured by the minimum description length (or bit-length) at which the observation is represented. For a distribution $P_{S,V}$, the optimal mapping of the data with parameter β , denoted by $P_{U|S}^{*,\beta}$, is found by solving the IB problem, defined as

$$L_{\beta}^{\text{IB},*}(P_{S,V}) := \max_{P_{U|S}} I(U; V) - \beta I(U; S), \quad (17)$$

over all mappings $P_{U|S}$ that satisfy $U - S - V$ and β is a positive Lagrange multiplier that allows to trade-off relevance and complexity. In Section VI-B several methods are discussed to obtain solutions $P_{U|S}^{*,\beta}$ to the IB problem in Eqn. (17) in several scenarios, e.g., when the distribution of (S, V) is perfectly known or only samples from it are available.

The IB problem is connected to multiple source-coding problems including source coding with logarithmic loss distortion [68], information combining [76], [77], common reconstruction [78], the

Wyner-Ahlsvede-Korner problem [79], [80], the efficiency of investment information [81]; to communications and cloud radio access networks (CRAN) [82], as well as learning including generalization [83], variational inference [84], representation learning and autoencoders [84], neural network compression [85], and others. See [86] and [87] for recent surveys on the IB principle and its application to learning. The connections between these problems allow extending results from one setup to another, and to consider generalizations of the classical IB problem to other setups including multi-terminal versions of the IB [88]–[90].

In fact, it is now well-known [91] that the IB problem in (17) is essentially a remote point-to-point lossy source-coding problem [92]–[94] where the distortion between the desired feature V^m and the reconstruction \hat{V}^m is measured under the logarithmic loss (log-loss) fidelity criterion [68]. That is, given m i.i.d. samples (V^m, S^m) , an encoder $f^{(m)}$ encodes the observation S^m using at most R bits per sample to generate an index $W = f^{(m)}(S^m)$. Using a decoder $g^{(m)}$, the receiver generates an estimate \hat{V}^m as a probability distribution over \mathcal{V}^m , i.e., $\hat{V}^m \in \mathcal{P}(V^m)$, such that $\hat{V}^m(V^m)$ is the value of that distribution evaluated at V^m given index w generated from S^m , i.e., $\hat{P}_{V^m|w}(V^m|w)$. The discrepancy between V^m and the estimation \hat{V}^m is measured using the m -letter distortion $d_{\log}(v^m, \hat{v}^m) := \frac{1}{m} \sum_{i=1}^m d_{\log}(v_i, \hat{v}_i)$, where

$$d_{\log}(v, \hat{v}) := \log \frac{1}{\hat{v}(v)}, \quad (18)$$

where $\hat{v}(v)$ is the value of that distribution evaluated at $v \in V$. The decoder $g^{(m)}$ is interested in a reconstruction to within an average distortion D , such that $\mathbb{E}[\frac{1}{m} \sum_{i=1}^m d_{\log}(v_i, \hat{v}_i)] \leq D$. The rate distortion function of this problem can be characterized as follows.

Theorem 4. (Remote Rate-Distortion Function with Log-Loss Distortion, [92], [94]) *The rate-distortion function for the remote source coding problem with log-loss distortion measure is given by*

$$R(D) = \min_{P_{U|S}: D \geq H(V|U)} I(U; S) \quad (19)$$

where the optimization is over all distributions satisfying $V - S - U$.

Using the substitution $\Delta := H(V) - D$, the region (R, D) described by this function can be seen to be equivalent to the convex hull of all pairs $(I(S; U), I(V; U))$ obtained from solving the IB problem in (17) for all β . Note that, for every $u \in U$ with $u = f(s)$ and $g(u) = Q(\cdot|u) \in P(V)$, we have

$$\begin{aligned} \mathbb{E}_{P_{S,V}}[d_{\log}(V, Q)] &= \sum_{s \in S, v \in V} P_{S,V}(s, v) \log \left(\frac{1}{Q(v|u)} \right) \\ &= H(V|S) + D_{KL}(P_{V|S} \| Q), \end{aligned} \quad (20)$$

which is minimized iff the estimate Q is given by the conditional posterior $P_V(V|U = u)$. Thus, operationally, the IB problem is equivalent to finding an encoder $P_{U|S}$ which maps the observation S to a representation U that satisfies the bit-rate constraint R , and such that U captures enough relevance of V so that the posterior probability of V given U minimizes the KL divergence between $P_{V|S}$ and the estimation $P_{V|U}$ produced by the decoder.

The IB is deeply linked to goal oriented compression, i.e., whenever compression is intended to perform a task. Consider the scenario in which a picture is taken at an edge device and a computational task, such as classifying an element in the picture or retrieving similar images to the one taken, needs to be performed at a remote unit. In classical compression, the image is compressed with the goal of preserving the maximum reconstruction quality before forwarding it over the communication channel. However, in goal-oriented compression only the information or features that are most relevant to perform the task need to be transmitted. For classification, those

features should be the most discriminative ones, and not necessarily a representation that allows to reconstruct the image, while for image retrieval, the task is different and so are the most relevant features. That is, the features are specific for the task, and the metric under which the task will be evaluated. This communication scenario is similar to that in Fig. 4, where S is a random variable modeling the observation, which is jointly distributed with the relevant information V for the task, which is not directly observable. The goal of the receiver is to recover an estimation of V with sufficient quality to perform the task and for the transmitter, to encode S into the representation U in order to send only the information necessary to recover V at the receiver, but not necessarily S . In the classical setup, in which the image needs to be reconstructed with minimum distortion we have $V = S$, while in classification, V can be the label class to classify the image to, such that the output obtained from the IB is the probability of the observation S belonging to the class given by V . How to select which are the relevant features for a given task is an open problem and depends on the metric that is used. However, in practise, often a careful encoding of the task into V and a conditional probability tailored for the task estimated by the IB to maximize the relevance can be a good candidate. This is justified, since the log-loss and the mutual information can be used to bound the performance of certain tasks, e.g., the probability of miss-classification of a classifier using a decision rule $Q_{\hat{V}|U}$ can be shown to be upper bounded as $\epsilon_{V|U}(Q_{\hat{V}|U}) := 1 - \mathbb{E}_{P_{S,V}}[Q_{\hat{V}|U}] \leq 1 - \exp \left(-\mathbb{E}_{P_{S,V}}[d_{\log}(V, Q_{\hat{V}|U})] \right)$ [83].

The formulation of the IB as a remote source coding under log-loss can be extended to consider context in the form of side information at the decoder or both the encoder and decoder [71] as described in Section III-B and studied in [95]. The IB problem can also be extended to multi-terminal scenarios. In particular, in [90], [96] the distributed oriented compression problem is studied from information-theoretic grounds using the IB formulation. In this scenario one is interested to perform a task at a remote destination, e.g., classification, on the reconstruction of the relevant information V reconstructed from the information relayed from K encoders, each observing some correlated information (S_1, \dots, S_K) with V . As shown in Fig. III-B, encoder k encodes the observation S_k into a representation U_k in order to preserve the most relevant information for the task. This problem can be shown to be essentially a K -encoder CEO problem under log-loss distortion [68], in which the decoder is interested in a soft estimate of V from the descriptions U_1, \dots, U_K generated from the observations S_1, \dots, S_K , each encoded with an average finite rate of R_k bits per sample. The fundamental limits in this scenario can be characterized in terms of optimal trade-off between relevance Δ and rate at each encoder,

Theorem 5. (Distributed Information Bottleneck Problem [90]) *The relevance-complexity region for the distributed IB problem is given by the union of all non-negative tuples $(\Delta, R_1, \dots, R_K)$ that satisfy*

$$\Delta \leq \sum_{k \in S} [R_k - I(S_k; U_k|V, T)] + I(V; S_{S^c}|T), \quad \forall S \subseteq K \quad (21)$$

for some distribution $p_T p_V \prod_{k=1}^K p_{S_k|V} \prod_{k=1}^K p_{S_k|S_k, T}$.

The IB framework can be used to design encoders and decoders in order to extract the most relevant information from observations. In Section VI-B we describe how to design encoders and decoders by solving (or approximating) the IB problem both for scenarios in which the distribution is perfectly known or when the source distribution is unknown and only samples are available.

More generally, the approach to extract the most relevant information for a task can also be considered for other end-to-end metrics that

go beyond log-loss by modelling the relevant information extraction as a remote source coding problem in which a metric needs to be minimized under some rate constraints, and the metric is a distortion (not necessarily additive) that represents the performance of the task, and can include, for example, the KL divergence, classification, hypothesis testing, etc. In other cases, the metric can also be learned implicitly by defining an alternative task, as in GANs [42], in which the generator and discriminator are trained to simultaneously extract the relevant information to generate realistic signals by learning how to pass an hypothesis test on distinguishing the generated data from real data.

E. Rate-Limited Remote Inference

As mentioned above, the hidden feature variable V can represent the class that the sample S belongs to, or the value associated to it in the case of a regression problem. If the goal is to convey only the class information to the receiver, and if there are no constraints on the complexity of the encoder, then the optimal operation would be to detect the class at the encoder, and only convey the class information to the receiver. From an information theoretic perspective, this problem can also be formulated as a remote inference problem.

In statistical inference problems, an observer observes n i.i.d. samples $S^n = (S_1, \dots, S_n)$ from some distribution p_S . We assume that this distribution belongs to a known family of distributions indexed by $V \in \mathcal{V}$; that is, S^n follows the conditional distribution $p_{S|V=v}(S^n) = \prod_{i=1}^n p_{S_i|V=v}(s_i)$. In general, the observer may want to estimate V simply from the available samples. If set \mathcal{V} is a discrete set, we have a detection/ *hypothesis testing* (HT) problem. If, instead, \mathcal{V} is a continuous set, i.e., $\mathcal{V} = \mathbb{R}^d$, then we have a *parameter estimation* problem. We impose a loss/ distortion function to quantify the quality of the estimation: $l(V, \hat{V})$, where \hat{V} is the estimate of the observer. The expected loss/risk of a decision rule $g^{(n)} : \mathcal{S}^n \rightarrow \mathcal{V}$, where $\hat{V} = g^{(n)}(S^n)$, is then defined as $R_v(g^{(n)}) = \mathbb{E}[l(V, \hat{V})]$.

In the Bayesian setting, we assume some known prior distribution on V , and try to minimize the average loss (also called *risk*) over the joint distribution of V and S , i.e., $\mathbb{E}[l(V, g^{(n)}(S^n))]$. Alternatively, we can also aim at minimizing the worst case loss/risk $R_{max} = \sup_{v \in \mathcal{V}} \mathbb{E}[l(v, \hat{V})]$. The corresponding decision rule is called the *minimax rule*. The classical Bayesian and minimax inference problems deal with centralized decision problems; that is, they assume the observer and the decision maker are the same agent, and makes the decision with full access to the samples. However, in many practical problems of interest, the observer and the decision maker are connected through a constrained communication channel. If we consider a rate-limited link, we obtain a remote inference problem. Remote inference problems over a rate-limited channel were first considered by Berger in [97]. We note that, the remote inference problem in the Bayesian setting is very similar to the information theoretic remote-rate distortion formulation in Section III-B. with the exception that we only have a single realization of the latent/hidden variable V , rather than a sequence of i.i.d. realizations each generating a separate sample S_i .

When the dimension of the parameter to be estimated is smaller than that of the observations, for example, in the case of HT, the observer can perform local inference and transmit its decision (indeed, optimal performance can be achieved asymptotically at *zero rate* by conveying the *type* information, which is a sufficient statistic [98]). Distributed parameter estimation with multiple terminals, each observing a component of a family of correlated samples, was studied in [99], under rate constraints from the observers to the decision maker in bits per sample. A single-letter bound, similar to the Shannon's rate distortion function, is provided on the variance of an

asymptotically unbiased estimator, which is later improved in [100]. Ahlswede and Burnashev studied the remote estimation problem when the decision maker has its own side information [101].

Ahlswede and Csiszár studied HT when the decision maker also has its own observations [102]. They studied the exponent of the type-II error probability, when a constrained is imposed on the type-I error probability. For the case of testing against independence, they were able to provide a single-letter expression similar to Shannon's rate-distortion function. This is one of the few cases in which a single-letter characterization is possible for a non-additive distortion measure. The more general distributed setting is considered in ???. Han showed in this paper that a positive exponent can be achieved even with a single-bit compression scheme. This result was extended to the more general zero-rate compression in [103]. This result was later refined by Shalaby and Papamarcou in [104], where they show that when the observers have fixed codebook sizes, the asymptotic performance does not depend on the particular codebook size. This means that no further gain can be obtained in terms of the asymptotic error exponents by allowing each observer to transmit a high-resolution soft decision instead of a binary decision. Despite significant research efforts, the optimal characterization of the type-II error exponent for the remote HT problem for the general case (beyond testing against independence) remains open to this day. Lower bounds were provided for the general problem in [102] and [105]. The distributed hypothesis testing is also studied in the context of a sensor network [106], where multiple sensors convey their noisy observations to a fusion center over rate-limited links. There has been a recent resurgence of interest in distributed HT problems [107]–[111].

IV. MACHINE LEARNING TECHNIQUES FOR SEMANTIC- AND TASK-ORIENTED COMPRESSION

The ultimate motivation of semantic compression is to extract the semantic information within the source data at the transmitter, which is most relevant for the task to be executed at the receiver. By filtering out task-irrelevant data both the bandwidth consumption and the transmission latency can be reduced significantly. However, the information theoretic framework presented above either assumes known statistics for the data and the relevant features for the task, or it is limited to the parameter estimation framework assuming i.i.d. samples from a family of distributions. On the other hand, in most practical applications we do not have access to statistical information, and often need to make inferences based on a single data sample. An alternative approach is to consider a data-driven framework, where we have access to a large dataset, which would allow us to train a model using machine learning tools to facilitate semantic information extraction without requiring a mathematical model. In particular, deep learning (DL) aided semantic extraction techniques have shown great potential for various information sources and associated tasks in the recent years.

While most machine learning research can be considered within the context of semantic feature extraction, we will focus on the communication aspects here. Machine learning algorithms typically follow a two-step approach: in the training phase, a model is trained using the available dataset for the desired task, e.g., classification or regression. Once the model is trained, it is used for prediction on new data samples. The communication in both phases can be considered in the context of semantic or goal-oriented communications. Below, we overview research in these two phases separately.

A. Remote Model Training

In the training phase, a single node or multiple nodes each with its own dataset communicate with a destination node with the goal

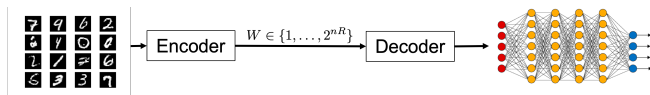


Fig. 6. Remote model training over a finite-rate communication link.

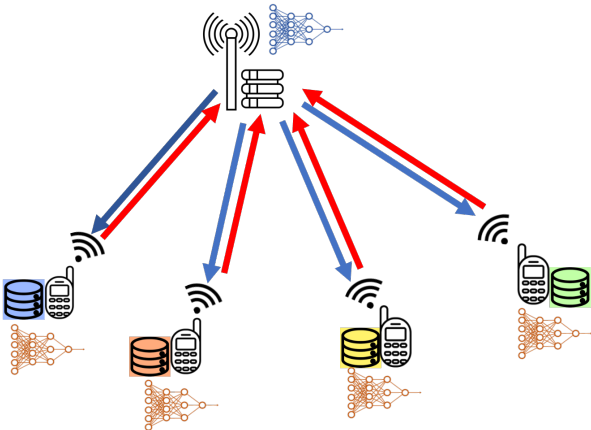


Fig. 7. Federated model training.

to reconstruct a model at the destination for a particular inference task. Note that this may also correspond to a storage problem, where the goal is to store the model in a limited storage to be later used in predicting future data samples. We would like to highlight here that this problem is an instance of a particular remote rate-distortion problem. Let us consider first the point-to-point version of the problem illustrated in Fig. 6. Here, we can treat the dataset at the encoder as the information available at the encoder, and the model itself as the remote source that the decoder wants to recover. Note that, similarly to the other semantic rate-distortion problems, the fidelity measure here is also not the similarity of the reconstructed neural network weights at the receiver to those trained at the encoder. In the end, what really matters is the performance of the reconstructed model at the decoder in terms of the prescribed quality measure, e.g., the accuracy of the reconstructed model at the decoder on new data samples.

As in the previous remote rate-distortion problems, a natural solution approach is to first estimate the remote source at the encoder; that is, to first train a model locally, and then convey this model to the decoder with the highest quality over the rate-limited channel, that is, in a way retaining the predictive power of the model on future queries. While the former step is simply the standard training process, the latter corresponds to model compression, which has been widely studied in recent years particularly in the context of deep neural networks (DNNs) that would otherwise require significant communication or storage resources.

There are various widely used methods to reduce the size of a pre-trained model. These include parameter pruning, quantization, low-rank factorization [112], and knowledge distillation. It has been known for a long time that many parameters in a neural network are redundant, and do not necessarily contribute significantly to the performance of the network. Therefore, redundant parameters that do not have a significant impact on the performance can be removed to reduce the network size and help address the overfitting problem [113]–[115]. In particular, the so-called ‘optimal brain surgeon’ in [115] uses the second-order derivative, i.e., the Hessian, of the loss function with respect to the network weights. In [114], the authors assumed that the Hessian matrix is diagonal, which causes the removal of incorrect weights. In [115], the authors showed that

the Hessian matrix is non-diagonal in most cases, and they proposed a more effective weight removal strategy. In addition to its better performance, the optimal brain surgeon does not require re-training after the pruning process. Note that pruning the network would also reduce the complexity and delay of the inference phase. Pruning is still a very active research area, and many different pruning methods have been studied, including weight, neuron, filter, or layer pruning. Please see [116] and references therein for a detailed survey on advanced pruning techniques. Another approach is to train the network directly with the sparsity constraints imposed during training, rather than reducing the network to a sparse one after training a full network [117], [118].

We can also apply quantization or other more advanced compression techniques, e.g., vector quantization, on the network parameters. Quantization has long been employed to reduce the network size for efficient storage [119], [120]. It is well-known that low precision representation of network weights is sufficient not just for inference based on trained networks but also for training them [121], [122]. At the extreme, it is possible to train DNNs even with single-bit binary weights [123], [124].

DNN compression can also be treated as a standard source compression problem, and vector quantization techniques can be employed for codebook-based compression to reduce the memory requirements. Similarly to pruning, Hessian-based quantization is shown to be effective in [125]. Hash functions are employed in [126] where the connections are hashed into groups, such that the ones in the same hash group share weights. It is argued in [127] that, for a typical network about 90% of the storage is taken up by the dense connected layers, while more than 90% of the running time is taken by the convolutional layers. Therefore, the authors focus on the compression of dense connected layers to reduce the storage and communication resources, and employ vector quantization to reduce the communication rate. In [128], Huffman coding is used to further compress the quantized network weights.

In knowledge transfer, the goal is to transfer the knowledge learned by a large, complex ensemble model into a smaller model without substantially reducing the network performance. It was first studied in [129]. In [130], the concept of temperature was introduced to generate the soft targets used for training the smaller model.

Another possible approach to solve this problem is model architecture optimization, where the goal is to adjust the size and complexity of DNN architectures to the constraints of the communication link during training without sacrificing their performance. Some of the popular recent efficient model architectures include SqueezeNet [131], MobileNets [132], ShuffleNet [133], and DenseNet [134]. We refer the reader to [135] for a more comprehensive survey of recent advances in model compression techniques for DNNs.

A more common scenario in remote model training is distributed training, in which multiple nodes each with its own local dataset collaborate to train a model by communicating with a remote parameter server, or with each other. The former scenario is known as *federated learning*, while the latter is referred to as *fully distributed*, or *peer-to-peer learning*. Please see Fig. 7 for an illustration of the federated learning scenario. Similarly to the single node scenario discussed above, federated training can be treated as a multi-terminal rate-distortion problem, where the datasets are observed samples at the multiple encoders, correlated with the underlying model, which is to be recovered at the parameter server. This would correspond to the CEO problem presented in Section III-B. Stochastic gradient descent based iterative algorithms are often used for federated learning. In the federated averaging (FedAvg) algorithm, proposed in [136], a global model is sent from the parameter server to the nodes, each of which computes a model update, typically employing multiple

stochastic gradient descent updates. The nodes then transmit these model updates back to the parameter server, which aggregates them, to finally update the global model. The algorithm is iterated until convergence. At each iteration of the algorithm, the goal is then to compute the average of the model updates, rather than the individual updates. Hence, this is a distributed lossy computation problem, which can be considered as yet another aspect of semantic communications. Here, the semantic that is relevant for the underlying task is a function of the multiple signals observed at different nodes.

Computation is often considered as a distinct problem from communication. One approach to computation over networks would be to carry out separate communication and computation steps. For example, if a node wants to compute a function of random variables S_1, \dots, S_n that are distributed over the network, we can first deliver these random variables to the node, which then computes the function value. In the point-to-point setting, the optimality of this approach can be shown in certain cases following the arguments of remote rate-distortion problem, where we treat the function to be computed as the latent variable V of the observed source S (see Fig. 4). However, in the general case, the optimal performance for a generic function is an open problem, even in the lossless computation case. The multi-terminal function computation problem was first introduced in [137], where the authors considered the parity function of two correlated symmetric binary random variables. They identified the optimal rate region for this case, and showed that this is not equivalent to the rate region one would obtain from [138] by first compressing and sending the observed sequences to the decoder. This illustrates the difficulty of the problem for arbitrary function computation. The problem was later studied in [139] in a point-to-point setting, considering one of the two sources is available at the decoder as side information. The optimal rate required for lossless computation of any function (in the Shannon theoretic sense - over long blocks with vanishing error probability) is characterized, and is shown to be given by the conditional G -entropy [140] of X given Y , where G is the characteristic graph of X, Y and function g to be computed as defined in [141]. While this is in general lower than first sending X to the decoder at rate $H(X|Y)$, and then computing the function, it is observed in [139] the gain is marginal in most cases.

B. Remote Inference

We next consider the machine learning approaches for rate-limited remote inference problems. Following the arguments in Section III, various lossy compression problems can be considered in the context of semantic communication under the appropriate reconstruction metric. In recent years, DNN aided compression algorithms have achieved significant results, often outperforming state-of-the-art standardised codecs in a variety of source domains, from image [?], [142]–[145], to video [146]–[154], speech [155], [156], and audio [157], [158] compression. One of the main advantages of DNN-based approaches compared to conventional compression algorithms is that they can be trained for any desired reconstruction metric at the receiver. For example, image or video compression algorithms can be trained with the SSIM or MS-SSIM metrics as objective, providing perceptually better reconstructions.

Task-oriented image compression was considered in [159], where the authors proposed lossy compression of MRI images to preserve as much clinically useful information as possible depending on the diagnostic task to be performed. In [160], the authors propose a metric based on conditional class entropy for a target detection task. In video compression, one approach is to employ region-of-interest compression, where only the relevant part of the video stream is compressed [161]–[163]. A classification aware distortion metric is

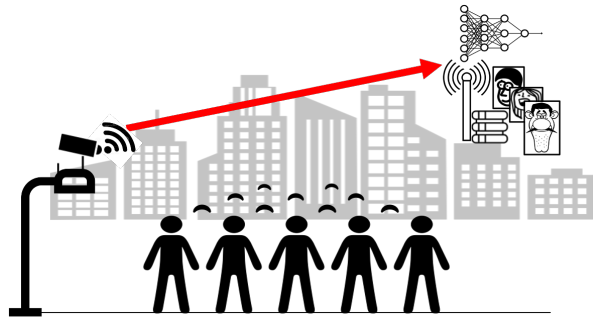


Fig. 8. Remote image retrieval problem.

proposed in [164], and applied to the High Efficiency Video Coding (HEVC) standard.

The authors of [165] have shown that latent representation produced by compressive autoencoders can be used to perform a classification task with ResNet-50 [166] network resulting in almost the same accuracy obtained by training on uncompressed image, showing that the classification network does not need to reconstruct the image first, at least explicitly. They also consider jointly training of the compression and classification networks. Task-based quantization is studied in [167] in the context of analog-to-digital conversion of signals for a specific task, and for channel estimation in [168].

We remark that the aforementioned task-based compression problems are remote rate-distortion problems in essence, but machine learning tools are employed mainly to acquire statistical knowledge from data. However, in these cases, since the transmitter has access to the original source information, the desired task can often be carried out at the transmitter, which is then transmitted to the receiver over the channel. Particularly for classification tasks this would require only a limited amount of information to be transmitted. A different type of remote inference from is considered in [169], called image retrieval. In this problem, illustrated in Fig. 8, the goal is to identify a query image of a person or a vehicle recorded locally by matching with images stored in a large database (gallery), typically available at the edge server. We emphasize that, unlike the typical classification tasks, the retrieval task cannot be performed locally as the database is available only at the remote edge server. The authors propose a retrieval-oriented image compression scheme, which compresses the feature vectors most relevant for the retrieval task, depending on the available bit budget.

V. SEMANTIC- AND TASK-ORIENTED COMMUNICATION OVER NOISY CHANNELS: JOINT SOURCE-CHANNEL CODING APPROACH

In the previous section, we have mainly focused on the compression aspects, assuming an error-free finite-rate communication channel from the encoder to the decoder. However, many communication channels suffer from noise, interference, and other imperfections. Shannon's channel coding theory mainly deals with communication over such noisy communication channels. However, as we have mentioned earlier, channel coding theory focuses on the reliable delivery of bits, whereas in the context of semantic communication, we will consider the transmission of source signals such as image, video, audio, or their features relevant for a particular task, over a noisy channel.

In this problem, illustrated in Figure V, the transmitter wants to transmit a sequence of source symbols $S^m \in \mathcal{S}^m$ over a noisy communication channel $P(Y^n|X^n)$, where $X^n \in \mathcal{X}^n$, $Y^n \in \mathcal{Y}^n$. Let $\hat{S}^m \in \hat{\mathcal{S}}^m$ denote the reconstruction of the receiver based on



Fig. 9. Illustration of a JSCC problem over a noisy communication channel.

Y^n . Similarly to the rate-distortion theory formulation, the goal is to minimize the distortion between S^m and \hat{S}^m under some given distortion (fidelity) measure, $d : \mathcal{S}^m \times \hat{\mathcal{S}}^m \rightarrow [0, \infty)$. More formally, let $f^{m,n} : \mathcal{S}^m \rightarrow \mathcal{X}^n$ denote the encoding function, and $g^{m,n} : \mathcal{Y}^n \rightarrow \hat{\mathcal{S}}^m$ denote the decoding function. In the case of an average distortion criteria, the goal is to identify the encoder and decoder function pairs which minimize $\mathbb{E}[d(S^m, \hat{S}^m)]$, where the expectation is over the source and channel distributions as well as any randomness the encoding and decoding functions may introduce. One can also impose an excess distortion criterion, where the goal is to minimize $\mathbb{P}[d(S^m, \hat{S}^m) > d]$, for some maximum allowable distortion target $d > 0$.

A (m, n) joint source-channel code of rate $r = m/n$ consists of an encoder-decoder pair, where the encoder $f^{(m,n)} : \mathcal{S}^m \rightarrow \mathcal{X}^n$ maps each source sequence s^m to a channel input sequence $x^n(s^m)$, and the decoder $g^{(m,n)} : \mathcal{Y}^n \rightarrow \hat{\mathcal{S}}^m$ maps the channel output y^n to an estimated source sequence \hat{s}^m . A rate-distortion pair (r, D) is said to be *achievable* if there exists a sequence of $(m, n(m))$ joint source-channel codes with rate r such that $n(m) \leq rm$, $\forall m$, and

$$\limsup_{m \rightarrow \infty} \mathbb{E}[d(S^m, \hat{S}^m(Y^n))] \leq D. \quad (22)$$

Shannon proved his famous Separation Theorem for a single-letter additive distortion measure; that is, $d(S^m, \hat{S}^m) = 1/m \sum_{i=1}^m d(S_i, \hat{S}_i)$ for the distortion measure $d(S, \hat{S}) \in [0, \infty)$. The theorem states the following.

Theorem 6. (Shannon's Separation Theorem, [6]) *Given a memoryless source S and a memoryless channel $p_{Y|X}$ with capacity $C = \sup_{p_X} p_X(x)I(X;Y)$, a rate-distortion pair (r, D) is achievable if $rR(D) < C$. Conversely, if a rate-distortion pair (r, D) is achievable, then $rR(D) \leq C$.*

The theorem states that we can separate the design of the communication system into two sub-problems without loss of optimality, the first focusing on the compression and the second focusing on the channel coding, each of them designed independently of the other. However, the optimality of separation holds only in the limit of infinite blocklength; whereas, in practice, it is possible to design joint source-channel codes that would outperform the best achievable separate code design. This was observed by Shannon in his 1959 paper [10]. Considering a binary source generating independent and equiprobable symbols and a memoryless binary symmetric channel, Shannon observed that simple uncoded transmission of symbols achieves the optimal distortion with rate $r = 1$ for one particular value of distortion D determined by the error probability of the channel. This observation was later extended by Goblick in [170] to Gaussian sources transmitted over Gaussian channels. This happens when the source distribution matches the optimal capacity achieving input distribution of the channel, and the channel at hand matches the optimal test channel achieving the optimal rate-distortion function of the source. Necessary and sufficient matching conditions are given in [171] for general source and channel distributions. However, these conditions are not satisfied for most practical source and channel distributions, and even when they hold, optimality of uncoded transmission fails when the coding rate is not 1, i.e., in the case of

bandwidth compression or expansion. On the other hand, the presence of such optimality results, that is, the fact that asymptotically optimal performance that requires infinite blocklength source and channel codes can be achieved by simple zero-delay uncoded transmission implies that there can be other non-separate coding schemes that can achieve near optimal performance, or outperform separation-based schemes in the finite blocklength regime.

A. Multi-terminal JSCC

It is well-known that the optimality of separation does not directly generalize to multi-terminal scenarios. This observation is often referred to the seminal work by Cover, El Gamal and Salehi [172], where they consider the transmission of correlated sources over a multiple access channel, and provide an example in which the uncoded transmission of the sources allow their perfect recovery, while this cannot be achieved by a separate scheme. Interestingly, a similar observation was already made by Shannon in [173] who considered the transmission of correlated sources over a two-way channel. The authors of [172] also proposed a coding technique (achievability result) exploiting the correlation among the sources. This result showed that instead of removing the correlation, we can utilize the dependency among the sources to design correlated channel codes, and in certain cases transmit the sources reliably even though this would not be possible with distributed compression followed by independent channel coding. Shortly afterwards a counter example was given in [174] showing that the sufficiency conditions provided in [172] are not necessary. More recently, [175] gave finite-letter sufficiency conditions. A similar problem of broadcasting correlated sources to multiple users is considered in [105], [176]. Many JSCC transmission strategies have been extensively studied for Gaussian sources over multi-terminal scenarios [177]–[181], and non-ergodic scenarios [182].

B. Remote Inference over Noisy Channels

In Subsection III-E, we have presented various statistical inference problems under rate constraints. We highlighted that these inference problems do not satisfy the additive single-letter requirement of typical Shannon theoretic distortion measures considered in the context of rate-distortion theory. Therefore, the separation theorem does not directly apply for these distortion measures.

Distributed hypothesis testing problem over a noisy communication channel was studied in [183] considering the type II error exponent (under a prescribed constraint on the type I probability of error) as the performance measure. Here, the task is to make a decision on the joint distribution of the samples observed by a remote observer and those observed by the decision maker. The observer communicates to the decision maker over a noisy channel. A separate hypothesis testing and channel coding scheme is presented, combining the Shimokawa-Han-Amari scheme [184] with a channel code that achieves the expurgated exponent with the best error-exponent for a single special message [185]. A joint scheme is also proposed using *hybrid coding* [186]. It is shown that separate scheme achieves the optimal type II error exponent when testing against independence. This is a special case of the problem, when the alternate hypothesis is the independence of the samples observed by the observer and the decision maker. While the optimal type II error exponent remains open in general, it is shown in [183] that joint encoding can strictly improve upon separation. This shows that communication and inference cannot be separated (without loss of optimality), but how the two should interact is vastly unexplored.

Distributed hypothesis testing over independent additive white Gaussian noise (AWGN) and fading channels, respectively, is studied

in [187] and [188]. These papers consider multiple sensors making noisy observations of the underlying hypothesis, and communicate to a fusion center over orthogonal noisy channels. Hypothesis testing over a discrete multiple access channel (MAC) is studied in [189], where the observations are quantized before being transmitted. Distributed estimation over a MAC is studied in [190], and a type-based uncoded transmission scheme is shown to be asymptotically optimal. Distributed estimation over a MAC is studied in [191] from a worst case risk point of view. Analog/uncoded transmission is again shown to outperform its digital separation-based counterparts, and to achieve a worst case risk that is within a logarithmic factor of an information theoretic lower bound.

VI. PRACTICAL DESIGNS FOR GOAL-ORIENTED COMMUNICATION OVER NOISY CHANNELS

Practical designs for JSCC of various information sources have been a long standing research challenge. Many different designs have been proposed in the literature, mainly based on the joint optimization of the parameters of an inherently separate design, i.e., channel dependent quantizer, [192]–[197]. Motivated by the theoretical optimality of uncoded transmission in certain ideal scenarios, analog transmission of discrete cosine transform (DCT) coefficients is proposed in [198] for wireless image transmission. In [199], the authors proposed linear coding of quantized wavelet coefficients. However, these efforts in JSCC design either do not provide sufficient performance gains or are too complex and specific to the underlying source and channel distributions to be applied in practice. Recently, JSCC schemes based on autoencoders [200], which are DNNs aimed at unsupervised data coding, have been introduced [201]–[205] and are shown to provide comparable or better performance than state-of-the-art separation-based digital schemes. In addition to improving the performance for a fixed channel state, these JSCC schemes also achieve graceful degradation with channel quality. That is, unlike separation-based approaches, their performance does not fall apart when the channel quality falls below a certain threshold. Similarly to DNN-aided data compression schemes, DNN-aided JSCC schemes also have the flexibility to adapt to a particular distortion measure, source domain, or a channel distribution.

The general schematic of a semantic communication system is shown in Fig. 10. The input data is sequentially passed through a semantic encoder and joint source-channel (JSC) encoder to extract semantic information relevant to receiver tasks, while receiver’s tasks could be either source data recovery or intelligent task execution. The benefits of this semantic communication approach is due to both the intelligent semantic encoding step, which basically extract the task-relevant features from the input, and the JSCC approach to delivering these features to the receiver.

According to the different types of source data, we will introduce the task-oriented semantic communications for text, speech, and image as follows.

1) *Task-Oriented Semantic Communications for Text*: For networks that allow interactions between humans as well as smart devices that have unique backgrounds and behavior patterns, reliable communication can be redefined as the *intended meaning* of messages being preserved at reception. Further, communicating parties can form social relationships and build trust, which may further affect how the received messages are interpreted. Motivated by these factors, reference [206] has proposed an approach that takes into account the meanings of the communicated messages and demonstrated the design of a point-to-point link to reliably communicate the meanings of messages through a noisy channel. To do so, in [206] the authors proposed a novel performance metric, *semantic error*, to measure how accurately the meanings of messages are recovered, and then

determined the optimal transmission policies to best preserve the meanings of recovered messages. This is achieved by leveraging lexical taxonomies and contextual information to design a graph-based index assignment scheme for fixed rate codeword assignment in a noisy channel. Words that are similar in their meaning (measured by the semantic distance between the words) are assigned to closer codewords in terms of their Hamming distance.

Building on this work, [207] modeled an external agent who can influence how the destination perceives the meaning of received information, to study the impact of social influence on contextual message interpretations on semantic communication. The exact nature of the agent, whether adversarial or helpful, is unknown to the communicating parties. [207] modeled this setup first as a Bayesian game played between the encoder/decoder and the influential entity. By extending the Bayesian game into a dynamic setting, the authors studied the interplay between the influential entity and the communicating parties, in which each player learned the true nature of the other player by updating its own beliefs as the game progresses, revealing information through observed actions. While these works pre-date the 6G efforts that brought semantics into the center stage, they serve as early efforts to build on designing semantic communication networks.

However, the semantic error of the aforementioned text-based semantic communication system is measured only at the word level instead of the sentence level. Thus, a transformer-powered semantic communication system for text, named DeepSC, was proposed in [208] by utilizing the meaning difference between the transmitted and received sentences. DeepSC can yield better performance than the traditional communication systems when coping with AWGN channels and it is more robust to channel variations, especially in the low SNR regime.

The core idea behind task-oriented semantic communications for text is to extract the useful information, e.g., grammatical information, word meanings, and logical relationships between words, to achieve intelligent tasks at the receiver, while ignoring the mathematical expression of words. In [209], Xie *et al.* have designed a multi-user semantic communication system to execute text-based tasks by transmitting text semantic features. Particularly, a transformer-enabled model, named DeepSC-MT, is proposed to perform the machine translation task for English-to-Chinese and Chinese-to-English by minimizing the meaning difference between sentences. The objective of DeepSC-MT is to map the meaning of the source sentence to the target language, which is achieved by learning the word distribution of the target language. Therefore, cross entropy is utilized as the loss function, represented as

$$\mathcal{L}_{\text{MT}} = \mathbb{E} \left[- \sum_n P(\mathbf{p}[n]) \log(P(\hat{\mathbf{p}}[n])) \right], \quad (23)$$

where $P(\mathbf{p}[n])$ and $P(\hat{\mathbf{p}}[n])$ are the real and predicted probability that the n -th word appears in the real translated sentence \mathbf{p} and the predicted translated sentence $\hat{\mathbf{p}}$, respectively.

Furthermore, another visual question answering (VQA) task has been investigated in [209] based on a multi-modal multi-user system. Particularly, the compressed text semantic features and image semantic features are extracted by a text semantic encoder and image semantic encoder at the transmitter, respectively, besides, a layer-wise transformer-enabled model is utilized at the receiver to perform the information query before fusing the image-text information to infer an accurate answer.

2) *Task-Oriented Semantic Communications for Speech*: The semantic extraction of speech signals is more complicated than text information. For speech signals, the semantic information required for transmission may refer to text information, emotional expression and type of language, etc., which increases the difficulty in

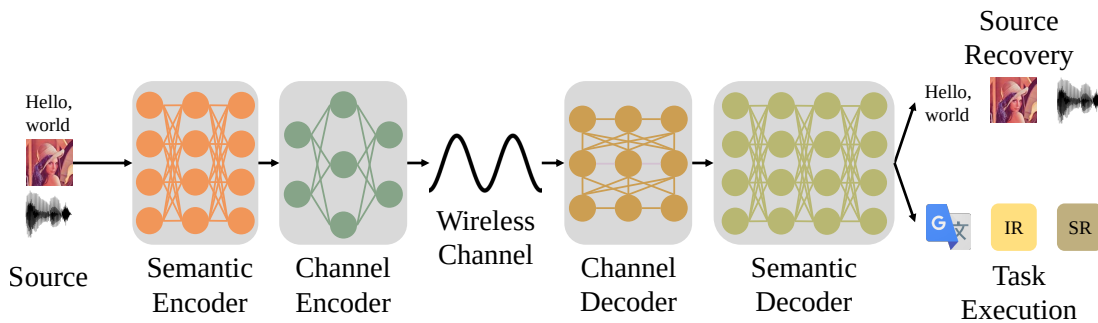


Fig. 10. Schematic of a DL-enabled semantic communication system.

extracting semantic features. Weng *et al.* has investigated a semantic communication system for speech signal reconstruction in [208], which aims to minimize the mean-square error (MSE) between the input and recovered speech sequences. Moreover, in [210], a speech recognition-oriented semantic communication system, named DeepSC-SR, has been developed to obtain the text transcription by transmitting the extracted text-related semantic features. Particularly, two convolutional layers are employed to constraint the input speech signals into low dimension representation before passing through the multiple gated recurrent unit (GRU)-based bidirectional recurrent neural network (BRNN) [211] modules. The text transcription is recognized at the character level by minimizing the difference of the character distribution between the source text sequence and the predicted text sequence. According to the connectionist temporal classification (CTC) [212], the loss function can be expressed as

$$\mathcal{L}_{CTC}(\theta) = -\ln \left(\sum_{A \in \mathfrak{A}(s, t)} p(\hat{t} | s, \theta) \right), \quad (24)$$

where $\mathfrak{A}(s, t)$ represents the set of all possible valid alignments of text sequence t to input speech s , $p(\hat{t} | s, \theta)$ denotes the posterior probability to recover one of the valid alignments \hat{t} based on s , and θ is the trainable parameters of the whole system.

Inspired by DeepSC-SR, a semantic communication system to achieve the speech recognition task at the word level has been proposed in [213], in which a visual geometry group (VGG)-enabled redundancy removal module is utilized to compress the transmitted data. The objective of this system is to convert the word distribution into the readable text transcription, which is achieved by the cross-entropy loss function between the source word sequence and predicted word sequence. Hyowoon *et al.* have developed a novel stochastic model of semantic-native communication (SNC) for generic tasks [214], which encourages the speaker to refer an entity by extracting the semantics and transmit its symbolic represent to a target listener. Besides, a further investigation is to fuse the local contextual reasoning into SNC, which allows the listener to obtain the most effective semantics from the listener at the transmitter. Furthermore, a curriculum learning framework for goal-oriented task execution has been investigated in [215], where the speaker describes the environment observations to enable the listener at the receiver to capture efficient semantics based on the defined language by using of the concepts of beliefs.

3) *Task-Oriented Semantic Communications for Image:* Furthermore, a semantic rate-distortion theory-based communication system for multiple image tasks has been investigated in [216], in which the source image is first restored at the receiver before intelligent task execution.

A semantic communication system for image retrieval over a wireless channel is considered in [169], which aims to identify the top- k similar images to a query image in a large dataset of images. The authors proposed directly mapping the extracted image features to channel inputs through a three-step training procedure: feature encoder pre-training, followed by JSCC autoencoder pre-training, and finally, end-to-end training. It is shown that this joint approach significantly outperforms the separation-based approach that combines the retrieval-oriented compression scheme mentioned in Section IV-B with a capacity-achieving channel code. This is despite the fact that a very short blocklength is considered, and hence, capacity is far from achievable. This scheme is recently extended in [209].

A semantic communication system for image classification has been proposed in [217], which performs a variational information bottleneck (VIB) framework to resolve the difficulty of mutual information computation of the original information bottleneck (IB) [75]. The adopted loss function can be expressed as

$$\mathcal{L}_{VIB}(\phi, \theta) = \mathbf{E}_{p(x, y)} \{ \mathbf{E}_{p_\phi(\hat{z}|x)} [-\log(q_\theta(y|\hat{z}))] + \beta D_{KL}(p_\phi(\hat{z}|x) \parallel q(\hat{z})) \}, \quad (25)$$

where x represents the input image, y denotes the target label, and \hat{z} is the recovered semantic information at the receiver. $q(\hat{z})$ and $q_\theta(y|\hat{z})$ are two variational distributions to approximate the true distributions of $p(\hat{z})$ and $p_\theta(y|\hat{z})$, respectively. ϕ and θ are the trainable parameters at the transmitter and receiver, respectively. $D_{KL}(\cdot)$ indicates the Kullback-Leibler divergence.

A. Distributed Training over Noisy Channels

We can also extend the model training tasks presented in Section IV-A for rate-limited channels to training over noisy channels. Since model training is often carried out over many iterations, training among wireless devices imposes strict delay constraints per iteration. Hence, the conventional approach of separate model compression and communication would not meet the desired delay and complexity requirements [218].

The problem of training and delivering a DNN to a remote terminal, called AirNet, is considered in [219] - extending the model in Fig. 6 by replacing the rate-limited error-free link with a noisy wireless channel. The conventional approach would be to first train a DNN, which is then compressed to be delivered efficiently over the bandwidth-limited channel. In the proposed delivery scheme in [219], the trained neural network weights are delivered over the wireless channel in an analog fashion; that is, they are mapped directly to the channel inputs. However, given the large size of DNNs, this would require a very large bandwidth. Moreover, the receiver will recover the noisy versions of the network weights. The authors propose two distinct strategies to remedy these problems.

Pruning [220] is employed to reduce the network size without sacrificing its performance significantly. The encoder first trains a large-dimensional DNN, which will be then pruned to the available channel bandwidth. Choosing a large DNN as an initial point, rather than directly training a DNN of dimension equal to the available channel bandwidth is motivated by the literature [221], which shows that pruning a trained large-dimensional DNN generally performs better than directly training a low-dimensional DNN. The noise problem is remedied by injecting a certain amount of noise to the network's weights at each training iteration, so that the trained network acquires robustness against channel noise. It is shown in [219] that the analog transmission of DNN weights achieves better accuracy compared to their digital delivery. A further unequal error protection strategy is also incorporated by pruning the network to a size smaller than the available channel bandwidth, and applying bandwidth expansion to a selected subset of more important weights using Shannon-Kotelnikov mapping [222].

While the above framework assumes the presence of the dataset at the encoder, in many practical settings, the encoder may simply have access to the DNN architecture and weights but no data. Delivering such a network over a wireless channel is considered in [223], where it is shown that a Bayesian approach at the receiver when estimating the noisy DNN weights can significantly improve its performance. The authors assume a Gaussian prior, and propose a population compensator and a bias compensator to the minimum mean square error (MMSE) metric.

For federated training at the edge (FEEL), uncoded and synchronized transmission of local model updates enables the parameter server to directly recover the sum of the updates from multiple terminals. This "over-the-air computation (OAC)" approach has received significant interest in recent years thanks to its bandwidth efficiency [224]–[227]. Instead of allocating orthogonal channel resources to the participating devices, they all share the same bandwidth. Unlike in conventional communication systems, the accuracy of the resultant computation benefits from more transmitters as the goal is to recover the sum of their model updates. OAC can also be used in a fully distributed learning scenario [228], where many computations take place in parallel.

B. Solutions to the Information Bottleneck Problem

The IB problem detailed in Section III-D provides a formulation to design mappings $P_{U|S}$ that allow to extract relevant information within the IB relevance–complexity region, i.e., the pairs of achievable (Δ, R) , by solving the IB problem in (17) for different values of β . However, in general, this optimization is challenging as it requires computation of mutual information terms. In this section, we describe how, for a fixed parameter β , the optimal solution $P_{U|S}^{\beta,*}$ or an efficient approximation of it, can be obtained under: (i) known general discrete memoryless distributions and particular distributions, or particular distributions such as Gaussian and binary symmetric sources; and (ii) unknown distributions and only samples are available to design the encoders and decoders.

1) *Known Discrete Memoryless Distribution*: When the relevant features V and the observation S is discrete and the joint distribution $P_{S,V}$ is known, the maximizing distributions $P_{U|S}$ in the IB problem in (17), can be efficiently found by an alternating optimization procedure similar to the expectation-maximization (EM) algorithm [229] and the standard Blahut–Arimoto (BA) method [230], [231], which is commonly used in the computation of rate-distortion functions of discrete memoryless sources. In particular, a solution $P_{U|X}$ to the constrained optimization problem is determined by the following self-consistent equations, for all $(u, s, v) \in U \times S \times V$, [75]

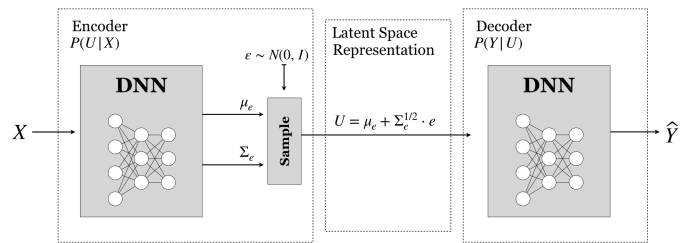


Fig. 11. Example parametrization of Variational Information Bottleneck using neural networks.

$$\begin{aligned}
 P_{U|S}(u|s) &= \frac{P_U(u)}{Z(\beta, s)} \exp\left(-\beta D_{\text{KL}}\left(P_{V|S}(\cdot|s) \| P_{V|U}(\cdot|u)\right)\right) \\
 P_U(u) &= \sum_{s \in S} P_S(s) P_{U|S}(u|s) \\
 P_{V|U}(v|u) &= \sum_{s \in S} P_{V|S}(v|s) P_{S|U}(s|u)
 \end{aligned} \tag{26}$$

where $P_{S|U}(s|u) = P_{U|S}(u|s) P_S(s) / P_U(u)$ and $Z(\beta, s)$ is a normalization term. It is shown in [75] that alternating iterations of these equations converges to a solution of the problem for any initial $P_{U|S}$. However, different to the standard Blahut–Arimoto algorithm for which convergence to the optimal solution is guaranteed, convergence here may be to a local optimum only.

2) *Unknown Distributions*: The main drawback of the solutions presented in the previous section is that requirement knowing the joint distribution $P_{S,V}$, or at least a good estimation of it and that iterating (26) can only be performed for sources with small alphabet (or jointly Gaussian [95], [232], [233]). The Variational IB (VIB) method was proposed in [84] as a means to obtain approximate solutions to the IB problem in the case in which the joint distribution is unknown and only a give training set of N samples $\{(s_i, v_i)\}_{i=1}^N$ is available or the alphabet is too large. The VIB consists in defining a variational (lower) bound on the cost of the IB problem in (17), use neural networks to parameterize this bound and show that, leveraging on the reparametrization trick [234] its optimization can be performed through stochastic gradient descent (SGD). From a task oriented communication perspective, the VIB approach provides a principled way to generalize the evidence lower bound (ELBO) and Variational Autoencoders [234] (and its extension to β -VAE cost [235]) to scenarios in which the decoder is interested in recovering the relevant information V for a task that is not necessarily the observed sample S by maximizing relevance. The idea is to use the IB principle to train an encoder and decoder, parameterized by DNNs, which are able to extract the relevant information to forward to a decoder in charge of reconstructing the relevant information. The resulting architecture to optimize with an encoder, a latent space, and a decoder parameterized by Gaussian distributions is shown in Figure 11. This approach has been used for task oriented communications also in JSCC scenarios, as a means to extract the relevant information to transmit over communication noisy channels to perform a given task at the destination, e.g., [236], [237].

More precisely, solving the IB problem in (17) consists in optimizing the IB-Lagrangian

$$L_{\beta}^{\text{IB}}(P_{U|S}) := I(U; V) - \beta I(U; S). \tag{27}$$

over all $P_{U|S}$ satisfying $U - S - V$. It follows from Gibbs inequality, for a any $P_{U|S}$ satisfying $U - S - V$, we have the following lower bound on the IB-Lagrangian,

$$L_{\beta}^{\text{IB}}(P_{U|S}) \geq L_{\beta}^{\text{VIB}}(P_{U|S}, Q_{V|U}, S_U) \tag{28}$$

$$:= \mathbb{E}_{P_{U|S}} [\log Q_{V|U}(V|U)] - \beta D_{\text{KL}}(P_{U|S} | S_U), \tag{29}$$

where $Q_{V|U}(V|u)$ is a given stochastic map $Q_{V|U} : U \rightarrow [0, 1]$ (also referred to as the variational approximation of $P_{V|U}$ or decoder) and $S_U(u) : U \rightarrow [0, 1]$ is a given stochastic map (also referred to as the variational approximation of P_U or prior), and $D_{\text{KL}}(P_{U|S}|S_U)$ is the relative entropy between $P_{U|S}$ and S_U . The equality holds iff $Q_{V|U} = P_{V|U}$ and $S_U = P_U$, i.e., the variational approximations match the true values.

Therefore, optimizing (27) over $P_{U|S}$ is equivalent to optimizing the variational cost (29) over $P_{U|S}$, $Q_{V|U}$ and S_U . In the VIB method, this optimization is done by further parameterizing the encoding and decoding distributions $P_{U|S}$, $Q_{V|U}$, and S_U that are to optimize using a family of distributions $P_\theta(u|s)$, $Q_\psi(v|u)$, and $S_\varphi(u)$, whose parameters are determined by DNNs with parameters θ, ϕ , and φ respectively. A common example is the family of multivariate Gaussian distributions [234], which are parameterized by the mean μ^θ and covariance matrix Σ^θ . Given an observation X , the values of $(\mu^\theta(s), \Sigma^\theta(s))$ are determined by the output of the DNN f_θ , whose input is S , and the corresponding family member is given by $P_\theta(u|s) = N(u; \mu^\theta(s), \Sigma^\theta(s))$. Another common example are Gumbel-Softmax distributions [238], [239].

The bound (29) restricted to family of distributions $P_\theta(u|s)$, $Q_\psi(v|u)$, and $S_\varphi(u)$ can be approximated using Monte-Carlo and the training samples $\{(s_i, v_i)\}_{i=1}^N$. To facilitate the computation of gradients using backpropagation [234], the reparameterization trick [234] is used to sample from $P_\theta(U|S)$. In particular, consider $P_\theta(U|S)$ to belong to a family of distributions that can be sampled by first sampling a random variable Z with distribution $P_Z(z)$, $z \in Z$ and then transforming the samples using some function $g_\theta : S \times Z \rightarrow U$ parameterized by θ , such that $U = g_\theta(s, Z) \sim P_\theta(U|s)$, e.g. a Gaussian distribution. The reparameterization trick is used to approximate by sampling M independent samples $\{u_m\}_{m=1}^M \sim P_\theta(u|s_i)$ for each training sample (s_i, v_i) , $i = 1, \dots, N$. Then, the lower bound on the IB cost can be optimized using optimization methods such as SGD or ADAM [240] with backpropagation over the the DNN parameters θ, ϕ, φ as,

$$\max_{\theta, \phi, \varphi} \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{\beta, i, M}^{\text{emp}}(\theta, \phi, \varphi), \quad (30)$$

where the cost for the i -th sample in the training dataset is:

$$\begin{aligned} & \mathcal{L}_{\beta, i, M}^{\text{emp}}(\theta, \phi, \varphi) \\ & := \frac{1}{M} \sum_{m=1}^M \left[\log Q_\phi(v_i|u_{i,m}) - \beta D_{\text{KL}}(P_\theta(U_i|s_i) \| Q_\varphi(U_i)) \right], \end{aligned} \quad (31)$$

and sampling is performed by using $u_{i,m} = g_\phi(s_i, z_m)$ with $\{z_m\}_{m=1}^M$ i.i.d. sampled from P_Z for each (s_i, v_i) pair.

For inference, let $\theta^*, \phi^*, \varphi^*$, be the DNN parameters obtained in training by solving (30). Inference for a new observation S , the representation U can be obtained by sampling from the encoders $P_{\theta_k^*}(U_k|S_k)$ and a soft estimate of the remote source Y can be inferred by sampling from the decoder $Q_{\phi^*}(V|U)$. Thus, from a task oriented communication perspective, $P_{\theta_k^*}(U_k|S_k)$ is an encoder trained according to the cost (30) to extract the most relevant representation U for inference of V and $Q_{\phi^*}(V|U)$ is a decoder that is trained to reconstruct the relevant information V from the representation U that minimized the log loss.

Similar to the steps followed for the variational IB in Section VI-B2 encoders and decoder performing on the region (21) can be computed by deriving a variational bound on and parameterizing encoding and decoding distributions $P_{U_k|S_k}, Q_{U_k|V}$ using a family of distributions whose parameters are determined by DNNs, and optimize it by using the reparameterization trick [234], Monte Carlo sampling, and stochastic gradient descent (SGD)-type algorithms. The encoders and

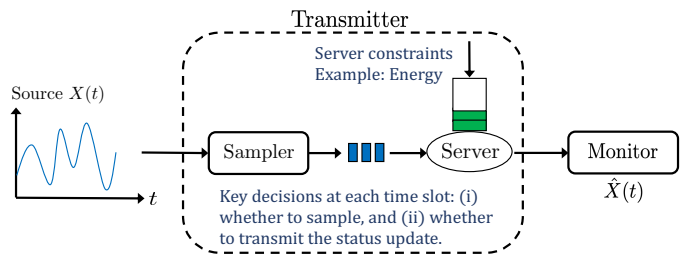


Fig. 12. An illustration of the general setup for Section VII.

decoders parameterized by DNN parameters θ, ϕ, φ , be optimized according to the distributed information bottleneck principles by considering the following empirical Monte Carlo approximation:

$$\begin{aligned} & \max_{\theta, \phi, \varphi} \frac{1}{n} \sum_{i=1}^n \left[\log Q_{\phi_K}(v_i|u_{1,i,j}, \dots, u_{K,i,j}) \right. \\ & \left. + s \sum_{k=1}^K \left(\log Q_{\phi_k}(v_i|u_{k,i,j}) - D_{\text{KL}}(P_{\theta_k}(U_{k,i}|s_{k,i}) \| Q_{\varphi_k}(U_{k,i})) \right) \right], \end{aligned} \quad (33)$$

where $u_{k,i,j} = g_{\phi_k}(s_{k,i}, z_{k,i,j})$ are samples obtained from the reparameterization trick by sampling from K random variables P_{Z_k} . The resulting architecture is shown in Figure 1. This architecture generalizes that from autoencoders to the distributed task oriented communication scenario setup with K encoders.

VII. TIMING AS SEMANTICS IN COMMUNICATIONS

This section deals with a broad class of problems where the relevance or value of information related to its timing. A specific sub-class of these problems relates to the popular idea of age of information (AoI). However, as we shall argue shortly, AoI is just one piece of the bigger puzzle when it comes to “timing as semantics” in communications. For a systematic exposition, we will first summarize key ideas related to the concept of AoI that will be useful in this discussion. We will then introduce a general real-time reconstruction problem using a rate-distortion viewpoint where semantic information is contained in the timing of the source samples. Along the way, we will explore the connection of age metrics to this rate-distortion viewpoint.

The general setup for this section is illustrated in Fig. 12. It consists of a source modeled as a random process $X(t)$, a transmitter that transmits status updates to a monitor about the current state of $X(t)$, and the monitor where the estimated state is denoted by $\hat{X}(t)$. In this setting, the transmitter needs to take two key decisions during each time slot: (i) whether to sample $X(t)$ to generate a status update, and (ii) whether to transmit the status update. Note that it is not always optimal to sample and transmit in the same time slot when the server is subjected to additional constraints, such as the energy availability/harvesting constraints. With this background, we now present some relevant background on AoI.

A. Age of Information

1) *Background and Definition*: AoI is a performance metric that quantifies freshness of information at a monitor about some remote stochastic process $X(t)$ observed by a transmitter node [241]–[244]. Specifically, AoI is defined as the time elapsed since the last successfully received update packet at the monitor was generated at the transmitter. Mathematically, it is defined as the following random process:

$$\Delta_{\text{AoI}}(t) = t - u(t), \quad (34)$$

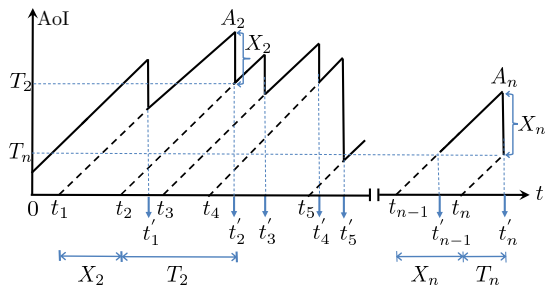


Fig. 13. AoI evolution vs. time for n update packets [243].

where $u(t)$ is the generation time instant of the latest status update received at the monitor by time t . In order to introduce the idea of AoI concretely, we use Fig. 13, which depicts a realization of AoI at the monitor as a function of time when the transmitter sends update packets according to a First-Come-First-Served (FCFS) queuing discipline and only one packet transmission may occur at any given time. Here, we implicitly model age as a linear function, which is also the most popular definition of age in the literature, although more general non-linear age functions have also been considered. Further, t_n and t'_n denote the generation and reception time instants of packet n at the transmitter and monitor, respectively. Therefore, we observe that: i) X_n is the inter-arrival time between packets $n-1$ and n , i.e., the time elapsed between the generation of packets $n-1$ and n , ii) T_n is the system time of packet n , i.e., the time elapsed from the generation of packet n at the transmitter until it is received at the monitor, and iii) AoI is reset to T_n at t'_n since packet n becomes the latest received update packet at t'_n , and hence the AoI value at that time instant is the time passed since the generation of packet n , which is T_n . Under ergodicity, many key properties of the AoI process, such as the average AoI or the mean peak AoI, can be studied from its sample functions, such as the one depicted in Fig. 13. Interested readers are advised to refer to well-known books and overview articles on this topic, such as [241]–[244], for a more comprehensive introduction to the concept of AoI.

2) *Age of Incorrect Information (AoII)*: While AoI is useful in characterizing the freshness of information and has found applications in a broad range of system settings under a variety of queuing disciplines [245]–[255], it does not explicitly take into account the similarity or discrepancy between the *status of information* at the transmitter and the monitor. In particular, as observed from (34), the AoI always increases as time passes even if the transmitter and monitor have the same status of information. As an example, consider $X(t)$ to be a discrete Markov source. If the current state of $X(t)$ matches with that of $\hat{X}(t)$, the monitor has the most updated information about $X(t)$ irrespective of when the status update was received. Therefore, in this case AoI is not an accurate measure of information freshness since it ignores semantic information about the nature of source $X(t)$. This motivated the introduction of a related metric, termed AoII, which is defined as [256]:

$$\Delta_{\text{AoII}} = (t - v(t)) \mathbf{1}\{X(t) \neq \hat{X}(t)\}, \quad (35)$$

where $v(t)$ is the last time instant when $\mathbf{1}\{X(t) \neq \hat{X}(t)\} = 0$ and $\mathbf{1}(\cdot)$ is an indicator function. Thus, according to (35), AoII grows with time only when $X(t) \neq \hat{X}(t)$. Therefore, AoII explicitly includes some semantic information about the nature of source $X(t)$ that was ignored in the definition of AoI. In order to make this connection more concrete, we now discuss the general real-time reconstruction problem from the rate-distortion viewpoint.

B. Real-time Reconstruction

We now establish a connection between the well-known age metrics discussed above and the semantics of information in the context of real-time remote tracking or reconstruction systems. Recall the setup of Fig. 12, where a transmitter observes source $X(t)$ and sends samples/updates about that source over time to a monitor. The objective is to reconstruct $X(t)$ at the monitor using the received samples/updates from the transmitter. Recall also that the transmitter has two key decisions to take during each time slot in this general setting: i) whether to sample $X(t)$ or not, and ii) whether to transmit the available or newly generated sample/update or not. Further, we consider a semantic-aware objective function (which is usually referred to as a *distortion function*), which can be expressed for the cases of discrete and continuous sources of information as

$$\mathcal{S}_{\text{dis}} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T c(t), \quad (36)$$

$$\mathcal{S}_{\text{con}} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T c(t) dt, \quad (37)$$

where $c(t)$ is an appropriately chosen cost function. For instance, $c(t)$ in the discrete case can be either $c(t) = |X(t) - \hat{X}(t)|$, $c(t) = a(t) \mathbf{1}\{X(t) \neq \hat{X}(t)\}$, or $c(t) = [X(t) - \hat{X}(t)]^2$, while the latter is also appropriate for the continuous case. The above problem can be thought of as a *rate-distortion problem*. The general objective here is to develop sampling and transmission policies that minimize the average distortion. Specific instances of this problem have already been investigated in the literature. For instance, when the transmitter samples and transmits during *every* time slot, the above setting reduces to the ones studied in [257] and [258]. Another specific example is when the transmitter samples every time slot and there is a deadline d for sending each sample, which corresponds to the setting studied in [259]. The focus of these analyses has been mainly on optimizing the transmission power of the transmitter in each time slot such that the distortion function is minimized.

Armed with this understanding, we now present three different sampling and transmission policies [15] that will also highlight subtle connections between the age metrics and semantics of information.

1) *AoI-aware Sampling and Transmission Policy*: In this policy, the transmitter samples a new update and sends it to the receiver once the AoI at the receiver reaches a predefined threshold A_{th} . While this policy takes into account the timeliness of the available information at the receiver, it is not semantics-aware since it does not involve the status of information at the source and the receiver in the process of decision-making. This again highlights the fact that AoI is agnostic of the source semantics.

2) *Source-aware Sampling and Transmission Policy*: In this policy, the transmitter samples a new update and sends it to the receiver when the status of information at the source changes. This policy can be thought of as a semantics-aware policy but just from the perspective of the source signal. Since this does not explicitly account for the discrepancy of the status at the transmitter and receiver, it could potentially result in unnecessary transmissions that could have been avoided in a truly semantics-aware policy. One specific example of such an instance given in [15] is when a transmission is triggered by a change in the status of the source but the status update does not reach the monitor because of an erasure. Now, if the source returns to the previous state in the next time slot, no transmission is required since there is no discrepancy in the states at the transmitter and receiver. However, the current policy will still transmit an update since it does not incorporate end-to-end semantics.

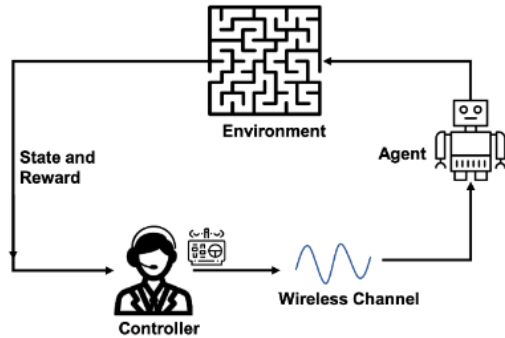


Fig. 14. Illustration of the effective communication paradigm introduced in [260].

3) *Semantics-aware Sampling and Transmission Policy*: Not surprisingly, this policy overcomes the drawback of the previous one by accounting for both the status of the source signal and the status of the reconstructed signal at the monitor. In particular, the transmitter samples a new update and sends it to the receiver only if the status of the source signal is different from that of the reconstructed signal. Here, we would like to remind the readers that this discrepancy on the status at the transmitter and the monitor explicitly appeared in the definition of AoII. Therefore, one can also interpret this policy as follows: the transmitter samples a new update and sends it to the receiver at the moment when the AoII becomes non-zero.

VIII. EFFECTIVE COMMUNICATIONS

In [5], Weaver defines Level C of communications as the effectiveness problem, which deals with the effect of the communicated symbols on the receiver. At this level, the communication should have not only a meaning, but also an impact on the behavior of the receiver. In [260], the authors argue that communication problems where the receiver takes actions based on the messages it receives, and the goal is to maximize a long-term reward function, falls into this category. The reward received following each transmission depends on the state of the receiver as well as the action it takes. In these problems, we cannot define a simple fidelity metric as in the preceding rate-distortion theory based framework. The utility, or fidelity of each transmission depends on its impact over a long time horizon. Here, as in the previous section, time plays an important role, and the same message may have different ‘meanings’ depending on the state of the receiver. Therefore, we can consider the action taken by the receiver after each transmission as the affect of communication on the receiver, and the state of the receiver as the ‘context’.

You can see in Fig. 14 an illustration of the effective communication paradigm proposed in [260]. The authors suggest that any single-agent Markov decision process can be turned into an effective communication problem by considering two entities, one of them, called the *controller*, observes the state and the reward signals, while the actions are taken by the *agent*, which may or may not observe its state, but not the rewards. We further assume that there is a noisy communication channel from the controller to the agent, and the actions taken by the agent will depend on the signals received from the controller through this channel. Note that the controller and the agent collaboratively need to learn the best actions to take to maximize their long term reward, which can be formulated as a reinforcement learning problem. However, unlike in conventional formulation, here the learning must be enabled through the noisy channel. It is shown in [260] that separating the

learning problem from the communication problem will result in a suboptimal performance. It is further shown in [261] that adjusting the communication scheme according to the state of the agent can result in a better long-term reward, providing evidence for context-dependent communications.

This formulation can be generalized to multiple agent scenarios, where the agents communicate with each other over noisy channels while trying to maximize prescribed reward functions [262]. This model generalizes also the recently popular models in the machine learning literature that study the *emergence of communication* among agents within the reinforcement learning literature [263]–[266].

IX. CONCLUSION

In this paper, we have presented a comprehensive overview of the information theoretical foundations of semantic- and task-oriented communications as well as practical data-driven approaches to semantic communication of various information sources including image, text and video. In particular, we argued that semantics can be considered in the context of rate-distortion theory, albeit for most practical information sources it is difficult if not impossible to explicitly state the corresponding distortion metric. Therefore, we provided an extensive overview of deep learning aided approaches to communicating practical information sources over constrained communication channels. In the case of rate-limited error-free communications, we have presented results related to remote inference as well as distributed model training, and exposed the relevance of information theoretic concepts and techniques such as information bottleneck. We emphasized the potential benefits of JSCC when communicating over noisy channels, both in terms of end-to-end performance and adaptivity to channel variations. Finally, we put the time dimension into the framework, and argued that the timeliness of a signal can be considered yet another aspect of semantics. Finally, we presented an interpretation of the effective communication problem, which involves communication among agents that take actions to achieve a certain goal. We believe that the provided in-depth overview of the information theoretical foundations of semantic- and task-oriented communications will provide guidelines for the development of more efficient practical solutions for semantic- and task-oriented communication networks.

REFERENCES

- [1] P. Popovski, O. Simeone, F. Boccardi, D. Gündüz, and O. Sahin, “Semantic-effectiveness filtering and control for post-5g wireless connectivity,” *Journal of Indian Inst. of Science*, 2020.
- [2] C. W. Morris, *Foundations of the Theory of Signs*. Chicago, IL: University of Chicago Press, 1938.
- [3] C. Cherry, *On Human Communication. A Review, a Survey, and a Criticism*. Cambridge, MA: John Wiley & Sons, 1957.
- [4] N. Chomsky, *Syntactic Structures*. The Hague/Paris: Mouton, 1957.
- [5] C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press, 1949.
- [6] C. E. Shannon, “A mathematical theory of communication,” *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [7] N. Wiener, *The Human Use of Human Beings*. Urbana, IL: Houghton Mifflin, 1950.
- [8] Y. Bar-Hillel and R. Carnap, “Semantic information,” *British Journal for the Philosophy of Science*, vol. 4, no. 14, pp. 147–157, 1953.
- [9] C. E. Shannon, “Prediction and entropy of printed english,” *The Bell System Technical Journal*, vol. 30, no. 1, pp. 50–64, 1951.
- [10] —, “Coding theorems for a discrete source with a fidelity criterion,” in *Institute of Radio Engineers, International Convention Record*, 1959.
- [11] Q. Lan, D. Wen, Z. Zhang, Q. Zeng, X. Chen, P. Popovski, and K. Huang, “What is semantic communication? a view on conveying meaning in the era of machine intelligence,” *J. Commun. Inform. Netw.*, vol. 6, no. 4, pp. 336–371, 2021.

- [12] E. C. Strinadi and S. Barbarossa, "6G networks: Beyond Shannon towards semantic and goal-oriented communications," *Computer Networks*, vol. 190, p. 107930, 2021.
- [13] M. Kalfa, M. Gok, A. Atalik, B. Tegin, T. M. Duman, and O. Arikan, "Towards goal-oriented semantic signal processing: Applications and future challenges," *Digital Signal Process.*, vol. 119, p. 103134, 2021.
- [14] G. Shi, Y. Xiao, Y. Li, and X. Xie, "From semantic communication to semantic-aware networking: Model, architecture, and open problems," *IEEE Commun. Mag.*, vol. 59, no. 8, pp. 44–50, 2021.
- [15] M. Kountouris and N. Pappas, "Semantics-empowered communication for networked intelligent systems," *IEEE Commun. Mag.*, vol. 59, no. 6, pp. 96–102, 2021.
- [16] S. Iyer, R. Khanai, D. Torse, R. J. Pandya, K. Rabie, K. Pai, W. U. Khan, and Z. Fadlullah, "A survey on semantic communications for intelligent wireless networks," *arXiv preprint arXiv:2202.03705*, 2022.
- [17] P. Zhang, W. Xu, H. Gao, K. Niu, X. Xu, X. Qin, C. Yuan, Z. Qin, H. Zhao, J. Wei *et al.*, "Toward wisdom-evolutionary and primitive-concise 6G: A new paradigm of semantic communication networks," *Engineering*, vol. 8, pp. 60–73, 2022.
- [18] Z. Qin, X. Tao, J. Lu, and G. Y. Li, "Semantic communications: Principles and challenges," *arXiv preprint arXiv:2201.01389*, 2021.
- [19] R. Carnap, Y. Bar-Hillel *et al.*, *An Outline of A Theory of Semantic Information*. RLE Technical Reports 247, Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge MA, Oct. 1952.
- [20] N. J. Venhuizen, M. W. Crocker, and H. Brouwer, "Semantic entropy in language comprehension," *Entropy*, vol. 21, no. 12, 2019.
- [21] I. D. Melamed, "Measuring semantic entropy," in *Tagging Text with Lexical Semantics: Why, What, and How?*, 1997.
- [22] X. Liu, W. Jia, W. Liu, and W. Pedrycz, "Afsse: An interpretable classifier with axiomatic fuzzy set and semantic entropy," *IEEE Trans. Fuzzy Syst.*, vol. 28, no. 11, pp. 2825–2840, 2020.
- [23] A. Chattopadhyay, B. D. Haeffe, D. Geman, and R. Vidal, "Quantifying task complexity through generalized information measures," <https://openreview.net/pdf?id=vcKVhY7AZqK>, 2020 (accessed 24th Dec. 2021).
- [24] Q. Wang, Z. Mao, B. Wang, and L. Guo, "Knowledge graph embedding: A survey of approaches and applications," *IEEE Trans. Knowledge and Data Eng.*, vol. 29, no. 12, pp. 2724–2743, Dec. 2017.
- [25] M. Nickel, V. Tresp, and H.-P. Kriegel, "A three-way model for collective learning on multi-relational data," in *Proc. 28th Int. Conf. Mach. Learn.*, 2011, pp. 809–816.
- [26] A. Bordes, X. Glorot, J. Weston, and Y. Bengio, "A semantic matching energy function for learning with multi-relational data," *Mach. Learn.*, vol. 94, no. 2, pp. 233–259, 2014.
- [27] Q. Liu, H. Jiang, A. Evdokimov, Z.-H. Ling, X. Zhu, S. Wei, and Y. Hu, "Probabilistic reasoning via deep learning: Neural association models," *arXiv preprint arXiv:1603.07704*, 2016.
- [28] S. Yang, J. Tian, H. Zhang, J. Yan, H. He, and Y. Jin, "Transms: Knowledge graph embedding for complex relations by multidirectional semantics," in *IJCAI*, 2019, pp. 1935–1942.
- [29] D. Collarana, M. Galkin, and I. Traverso-Rib³n.
- [30] K. Zhou, W. X. Zhao, S. Bian, Y. Zhou, J.-R. Wen, and J. Yu, "Improving conversational recommender systems via knowledge graph based semantic fusion," in *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1006–1014.
- [31] T. Grainger, K. Aljadda, M. Korayem, and A. Smith, "The semantic knowledge graph: A compact, auto-generated model for real-time traversal and ranking of any relationship within a domain," in *IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, 2016, pp. 420–429.
- [32] Z. Liu, C. Xiong, M. Sun, and Z. Liu, "Entity-duet neural ranking: Understanding the role of knowledge graph semantics in neural information retrieval," *arXiv preprint arXiv:1805.07591*, 2018.
- [33] C. Xiong, R. Power, and J. Callan, "Explicit semantic ranking for academic search via knowledge graph embedding," in *Proceedings of the 26th international conference on world wide web*, 2017, pp. 1271–1279.
- [34] S. Purohit, N. Van, and G. Chin, "Semantic property graph for scalable knowledge graph analytics," in *IEEE International Conference on Big Data (Big Data)*, 2021, pp. 2672–2677.
- [35] S. Bhatt, A. Sheth, V. Shalin, and J. Zhao, "Knowledge graph semantic enhancement of input data for improving AI," *IEEE Internet Computing*, vol. 24, no. 2, pp. 66–72, Mar. 2020.
- [36] Z. Qi, Z. Zhang, J. Chen, X. Chen, Y. Xiang, N. Zhang, and Y. Zheng, "Unsupervised knowledge graph alignment by probabilistic reasoning and semantic embedding," *arXiv preprint arXiv:2105.05596*, 2021.
- [37] H. Xiao, M. Huang, and X. Zhu, "Knowledge semantic representation: a generative model for interpretable knowledge graph embedding," *arXiv preprint arXiv:1608.07685*, 2016.
- [38] F. Zhou, Y. Li, X. Zhang, Q. Wu, X. Lei, and R. Q. Hu, "Cognitive semantic communication systems driven by knowledge graph," *arXiv preprint arXiv:2202.11958*, 2022.
- [39] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *The Thirty-Seventh Asilomar Conference on Signals, Systems Computers, 2003*, vol. 2, Nov. 2003, pp. 1398–1402 Vol.2.
- [40] Y. Blau and T. Michaeli, "Rethinking lossy compression: The rate-distortion-perception tradeoff," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 675–685.
- [41] Z. Wang, G. Wu, H. Sheikh, E. Simoncelli, E.-H. Yang, and A. Bovik, "Quality-aware images," *IEEE Transactions on Image Processing*, vol. 15, no. 6, pp. 1680–1689, 2006.
- [42] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014. [Online]. Available: <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>
- [43] M. Tschannen, E. Agustsson, and M. Lucic, "Deep generative models for distribution-preserving lossy compression," in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018. [Online]. Available: <https://proceedings.neurips.cc/paper/2018/file/801fd8c2a4e79c1d24a40dc735c051ae-Paper.pdf>
- [44] E. Agustsson, M. Tschannen, F. Mentzer, R. Timofte, and L. V. Gool, "Generative adversarial networks for extreme learned image compression," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [45] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 06–11 Aug 2017, pp. 214–223.
- [46] S. Nowozin, B. Cseke, and R. Tomioka, "f-gan: Training generative neural samplers using variational divergence minimization," 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. [Online]. Available: <https://proceedings.neurips.cc/paper/2016/file/cedebb6e872f539bef8c3f919874e9d7-Paper.pdf>
- [47] M. Akbari, J. Liang, and J. Han, "Dsslic: Deep semantic segmentation-based layered image compression," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 2042–2046.
- [48] K. Mahalingaiah, H. Sharma, P. Kaplish, and I. Cheng, "Semantic learning for image compression (slc)," in *Smart Multimedia*, T. McDaniel, S. Berretti, I. D. D. Curcio, and A. Basu, Eds. Cham: Springer International Publishing, 2020, pp. 57–66.
- [49] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [50] J. Fu, J. Liu, H. Tian, Y. Li, Y. Bao, Z. Fang, and H. Lu, "Dual attention network for scene segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [51] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *CoRR*, vol. abs/1411.1784, 2014. [Online]. Available: <http://arxiv.org/abs/1411.1784>
- [52] J. Johnson, R. Krishna, M. Stark, L.-J. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei, "Image retrieval using scene graphs," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3668–3678.
- [53] X. Chang, P. Ren, P. Xu, Z. Li, X. Chen, and A. G. Hauptmann, "A comprehensive survey of scene graphs: Generation and application," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.

- [54] R. Wang, Z. Wei, P. Li, Q. Zhang, and X. Huang, "Storytelling from an image stream using scene graphs," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, pp. 9185–9192, Apr. 2020. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/6455>
- [55] P. Manoranjan, M. Murshed, and L. S. Dooley, "Very low bit-rate video coding," in *Video Data Management and Information Retrieval*, S. Deb, Ed. Hershey, PA: IGI Global, 2005, pp. 100–148.
- [56] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, no. 1, pp. 185–203, 1981.
- [57] J. Ostermann, "Object-based analysis-synthesis coding based on the source model of moving rigid 3d objects," *Signal Processing: Image Communication*, vol. 6, no. 2, pp. 143–161, 1994. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0923596594900124>
- [58] R. Talluri, K. Oehler, T. Barmon, J. Courtney, A. Das, and J. Liao, "A robust, scalable, object-based video compression technique for very low bit-rate coding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 7, no. 1, pp. 221–233, 1997.
- [59] L. Duan, J. Liu, W. Yang, T. Huang, and W. Gao, "Video coding for machines: A paradigm of collaborative compression and intelligent analytics," *Trans. Image Processing*, vol. 29, pp. 8680–8695, jan 2020.
- [60] R. Gbadeyan and C. Joslin, "Object based hybrid video compression," in *Proce. Int'l Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 4: VISAPP, INSTICC*. SciTePress, 2021, pp. 785–792.
- [61] R. Dobrushin and B. Tsybakov, "Information transmission with additional noise," *IRE Transactions on Information Theory*, vol. 8, no. 5, pp. 293–304, 1962.
- [62] J. Liu, W. Zhang, and H. V. Poor, "A rate-distortion framework for characterizing semantic information," in *IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 2894–2899.
- [63] P. A. Stavrou and M. Kountouris, "A Rate Distortion Approach to Goal-Oriented Communication," *techrxiv*, 2 2022. [Online]. Available: https://www.techrxiv.org/articles/preprint/A_Rate_Distortion_Approach_to_Goal-Oriented_Communication/19128026
- [64] J. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion," *IEEE Transactions on Information Theory*, vol. 16, no. 4, pp. 406–411, 1970.
- [65] T. Berger, Z. Zhang, and H. Viswanathan, "The ceo problem [multiterminal source coding]," *IEEE Transactions on Information Theory*, vol. 42, no. 3, pp. 887–902, 1996.
- [66] H. Viswanathan and T. Berger, "The quadratic Gaussian CEO problem," *IEEE Transactions on Information Theory*, vol. 43, no. 5, pp. 1549–1559, 1997.
- [67] Y. Oohama, "Rate-distortion theory for Gaussian multiterminal source coding systems with several side informations at the decoder," *IEEE Transactions on Information Theory*, vol. 51, no. 7, pp. 2577–2593, 2005.
- [68] T.-A. Courtade and T. Weissman, "Multiterminal source coding under logarithmic loss," *IEEE Trans. on Info. Theory*, vol. 60, pp. 740–761, Jan. 2014.
- [69] Y. Ugur, I.-E. Aguerri, and A. Zaidi, "Vector gaussian ceo problem under logarithmic loss," *IEEE Trans. on Info. Theory*. Submitted for publication, 2018.
- [70] R. M. Gray, "Conditional rate-distortion theory," pp. 6502–2, Oct 1972.
- [71] A. D. Wyner and J. Ziv, "The rate-distortion function for source coding with side information at the decoder," *IEEE Trans. Inf. Theory*, vol. 22, pp. 1–10, Jan. 1976.
- [72] A. Aaron, E. Setton, and B. Girod, "Towards practical wyner-ziv coding of video," in *Proceedings 2003 International Conference on Image Processing (Cat. No.03CH37429)*, vol. 3, 2003, pp. III–869.
- [73] A. Aaron, R. Zhang, and B. Girod, "Wyner-ziv coding of motion video," in *Conference Record of the Thirty-Sixth Asilomar Conference on Signals, Systems and Computers, 2002.*, vol. 1. IEEE, 2002, pp. 240–244.
- [74] R. Gray and A. Wyner, "Source coding for a simple network," *Bell System Technical Journal*, vol. 53, no. 9, pp. 1681–1721, 1974.
- [75] N. Tishby, F. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annu. Allerton Conf. Commun., Control, Comput.*, 1999, pp. 368–377.
- [76] I. Sutskever, S. Shamai, and J. Ziv, "Extremes of information combining," *IEEE Trans. Inform. Theory*, vol. 51, no. 04, pp. 1313–1325, 2005.
- [77] I. Land and J. Huber, "Information combining," *Foundations and Trends in Commun. and Inform. Theory*, vol. 03, pp. 227–230, Nov. 2006.
- [78] Y. Steinberg, "Coding and common reconstruction," *IEEE Trans. Inf. Theory*, vol. IT-11, pp. 4995–5010, Nov. 2009.
- [79] A. D. Wyner, "On source coding with side information at the decoder," vol. IT-21, pp. 294–300, May 1975.
- [80] R. Ahlswede and J. Korner, "Source coding with side information and a converse for degraded broadcast channels," *IEEE Trans. Inf. Theory*, vol. 21, no. 6, pp. 629–637, November 1975.
- [81] E. Erkip and T. M. Cover, "The efficiency of investment information," *IEEE Trans. Info. Theory*, vol. 44, no. 3, pp. 1026–1040, May 1998.
- [82] I. Estella Aguerri, A. Zaidi, G. Caire, and S. Shamai Shitz, "On the capacity of cloud radio access networks with oblivious relaying," *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4575–4596, 2019.
- [83] M. Vera, P. Piantanida, and L. R. Vega, "The role of information complexity and randomization in representation learning." [Online]. Available: <https://arxiv.org/abs/1802.05355>
- [84] A. Alemi, I. Fischer, J. Dillon, and K. Murphy, "Deep variational information bottleneck," in *ICLR*, 2017. [Online]. Available: <https://arxiv.org/abs/1612.00410>
- [85] B. Dai, C. Zhu, B. Guo, and D. Wipf, "Compressing neural networks using the variational information bottleneck," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1135–1144.
- [86] A. Zaidi, I. Estella-Aguerri, and S. Shamai, "On the information bottleneck problems: Models, connections, applications and information theoretic views," *Entropy*, vol. 22, no. 2, p. 151, 2020.
- [87] Z. Goldfeld and Y. Polyanskiy, "The information bottleneck problem and its applications in machine learning," *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 19–38, 2020.
- [88] Q. Yang, P. Piantanida, and D. Gündüz, "The multi-layer information bottleneck problem," in *2017 IEEE Information Theory Workshop (ITW)*, 2017, pp. 404–408.
- [89] M. Vera, L. Rey Vega, and P. Piantanida, "Collaborative information bottleneck," *IEEE Transactions on Information Theory*, vol. 65, no. 2, pp. 787–815, 2019.
- [90] I. E. Aguerri and A. Zaidi, "Distributed variational representation learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 1, pp. 120–138, 2021.
- [91] P. Haremoes and N. Tishby, "The information bottleneck revisited or how to choose a good distortion measure," in *Proc. IEEE Int. Symp. Information Theory*, Jun. 2007, pp. 566–570.
- [92] R.-L. Dobrushin and B. S. Tsybakov, "Information transmission with additional noise," *IRE Trans. on Info. Theory*, vol. 85, pp. 293–304, 1962.
- [93] H.-S. Witsenhausen and A.-D. Wyner, "A conditional entropy bound for a pair of discrete random variables," *IEEE Trans. on Info. Theory*, vol. IT-21, pp. 493–501, Sep. 1975.
- [94] H.-S. Witsenhausen, "Indirect rate distortion problems," *IEEE Trans. on Info. Theory*, vol. IT-26, pp. 518–521, Sep. 1980.
- [95] Y. Uğur, I. Estella-Aguerri, and A. Zaidi, "Vector gaussian ceo problem under logarithmic loss and applications," *IEEE Transactions on Information Theory*, vol. 66, no. 7, pp. 4183–4202, 2020.
- [96] I. E. Aguerri and A. Zaidi, "Distributed information bottleneck method for discrete and Gaussian sources," in *Proc. of Int. Zurich Seminar on Information and Communication, IZS*, Zurich, Switzerland, 2018.
- [97] T. Berger, "Decentralized estimation and decision theory," in *IEEE 7th. Spring Workshop on Inf. Theory*, Mt. Kisco, NY, Sep. 1979.
- [98] T. S. Han and S. Amari, "Statistical inference under multiterminal data compression," *IEEE Trans. Information Theory*, vol. 44, no. 6, pp. 2300–2324, Oct. 1998.
- [99] Z. Zhang and T. Berger, "Estimation via compressed information," *IEEE Trans. on Info. Theory*, vol. 34, no. 2, pp. 198–211, March 1988.
- [100] T. S. Han and S. Amari, "Parameter estimation with multiterminal data compression," *IEEE Trans. on Info. Theory*, vol. 41, no. 6, pp. 1802–1833, Nov 1995.
- [101] R. Ahlswede and M. V. Burnashev, "On minimax estimation in the presence of side information about remote data," *The Annals of Statistics*, vol. 18, no. 1, p. 141â71, 1990.
- [102] R. Ahlswede and I. Csiszár, "Hypothesis testing with communication constraints," *IEEE Trans. Information Theory*, vol. 32, no. 4, pp. 533–542, Jul. 1986.
- [103] T. Han and K. Kobayashi, "Exponential-type error probabilities for multiterminal hypothesis testing," *IEEE Transactions on Information Theory*, vol. 35, no. 1, pp. 2–14, 1989.
- [104] H. Shalaby and A. Papamarcou, "Multiterminal detection with zero-rate data compression," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 254–267, 1992.
- [105] T. Han, "Hypothesis testing with multiterminal data compression," *IEEE Transactions on Information Theory*, vol. 33, no. 6, pp. 759–772, 1987.

- [106] J.-F. Chamberland and V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 407–416, 2003.
- [107] M. S. Rahman and A. B. Wagner, "On the optimality of binning for distributed hypothesis testing," *IEEE Transactions on Information Theory*, vol. 58, no. 10, pp. 6282–6303, 2012.
- [108] G. Katz, P. Piantanida, and M. Debbah, "Distributed binary detection with lossy data compression," *IEEE Trans. Information Theory*, vol. 63, no. 8, pp. 5207–5227, Mar. 2017.
- [109] S. Salehkalaibar, M. Wigger, and L. Wang, "Hypothesis testing over the two-hop relay network," *IEEE Trans. Information Theory*, vol. 65, no. 7, pp. 4411–4433, Jul. 2019.
- [110] N. Weinberger and Y. Kochman, "On the reliability function of distributed hypothesis testing under optimal detection," *IEEE Trans. Information Theory*, vol. 65, no. 8, pp. 4940–4965, Aug. 2019.
- [111] P. Escamilla, M. Wigger, and A. Zaidi, "Distributed hypothesis testing: Cooperation and concurrent detection," *IEEE Transactions on Information Theory*, vol. 66, no. 12, pp. 7550–7564, 2020.
- [112] E. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus, "Exploiting linear structure within convolutional networks for efficient evaluation," in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS'14. Cambridge, MA, USA: MIT Press, 2014, p. 1269a1277.
- [113] S. Hanson and L. Pratt, "Comparing biases for minimal network construction with back-propagation," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, Denver, CO, USA, Nov. 1989.
- [114] Y. LeCun, J. Denker, and S. Solla, "Optimal brain damage," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, Denver, CO, USA, Nov. 1990.
- [115] B. Hassibi, D. G. Stork, G. Wolff, and T. Watanabe, "Optimal brain surgeon: Extensions and performance comparisons," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, San Francisco, CA, USA, Jul. 1993.
- [116] J. Liu, S. Tripathi, U. Kurup, and M. Shah, "Pruning algorithms to accelerate convolutional neural networks for edge applications: A survey," *arXiv preprint arXiv:2005.04275*, 2020.
- [117] V. Lebedev and V. Lempitsky, "Fast ConvNets using group-wise brain damage," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, June 2016.
- [118] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, Barcelona, Spain, Dec. 2016.
- [119] C. H. Fiesler E, Choudry A, "Weight discretization paradigm for optical neural networks," in *Optical Interconnections and Networks, Proc. of International Society for Optics and Photonics*, vol. 1281, 1990, p. 164a174.
- [120] W. Balzer, M. Takahashi, J. Ohta, and K. Kyuma, "Weight quantization in boltzmann machines," *Neural Networks*, vol. 4, no. 3, pp. 405–409, 1991. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0893608091900771>
- [121] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, "Deep learning with limited numerical precision," in *Proc. International Conference on Machine Learning (ICML)*, Lille, France, Jul. 2015.
- [122] M. Courbariaux, Y. Bengio, and J. P. David, "Training deep neural networks with low precision multiplications," in *Proc. International Conference on Learning Representations Workshop (ICLR Workshop)*, San Diego, CA, USA, May. 2015.
- [123] —, "Binaryconnect: Training deep neural networks with binary weights during propagations," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, Montreal Canada, Dec. 2015.
- [124] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1," *arXiv preprint arXiv:1602.02830*, 2016.
- [125] Y. Choi, M. El-Khamy, and J. Lee, "Towards the limit of network quantization," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017. [Online]. Available: <https://openreview.net/forum?id=rJ8uNptgl>
- [126] W. Chen, J. T. Wilson, S. Tyree, K. Q. Weinberger, and Y. Chen, "Compressing neural networks with the hashing trick," in *Proc. International Conference on Machine Learning (ICML)*, Lille, France, Jul. 2015.
- [127] Y. Gong, L. Liu, M. Yang, and L. Bourdev, "Compressing deep convolutional networks using vector quantization," *arXiv preprint arXiv:1412.6115*, 2014.
- [128] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," in *Proc. International Conference on Learning Representations (ICLR)*, San Juan, Puerto Rico, May. 2016.
- [129] C. Bucilau, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 535a541. [Online]. Available: <https://doi.org/10.1145/1150402.1150464>
- [130] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, cite arxiv:1503.02531. [Online]. Available: <http://arxiv.org/abs/1503.02531>
- [131] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 10x smaller model size," 2016. [Online]. Available: <https://arxiv.org/abs/1602.07360>
- [132] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," 2017. [Online]. Available: <http://arxiv.org/abs/1704.04861>
- [133] X. Zhang, X. Zhou, M. Lin, and J. Sun, "Shufflenet: An extremely efficient convolutional neural network for mobile devices," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6848–6856.
- [134] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2261–2269.
- [135] L. Deng, G. Li, S. Han, L. Shi, and Y. Xie, "Model compression and hardware acceleration for neural networks: A comprehensive survey," *Proceedings of the IEEE*, vol. 108, no. 4, pp. 485–532, 2020.
- [136] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Singh and J. Zhu, Eds., vol. 54. PMLR, 20–22 Apr 2017, pp. 1273–1282. [Online]. Available: <https://proceedings.mlr.press/v54/mcmahan17a.html>
- [137] J. Korner and K. Marton, "How to encode the modulo-two sum of binary sources (corresp.)," *IEEE Transactions on Information Theory*, vol. 25, no. 2, pp. 219–221, Mar 1979.
- [138] T. Berger, *Multiterminal source coding*. New York, NY, USA: Springer-Verlag, 1977.
- [139] A. Orlitsky and J. Roche, "Coding for computing," *IEEE Transactions on Information Theory*, vol. 47, no. 3, pp. 903–917, 2001.
- [140] G. Simonyi, "Graph entropy: A survey," *Combinatorial Optimization*, vol. 20, pp. 399–441, 1995.
- [141] H. Witsenhausen, "The zero-error side information problem and chromatic numbers (corresp.)," *IEEE Transactions on Information Theory*, vol. 22, no. 5, pp. 592–593, 1976.
- [142] G. Toderici, S. M. O'Malley, S. J. Hwang, D. Vincent, D. Minnen, S. Baluja, M. Covell, and R. Sukthankar, "Variable rate image compression with recurrent neural networks," *Proc. of the Int. Conf. on Learning Representations (ICLR)*, 2016.
- [143] O. Rippel and L. Bourdev, "Real-time adaptive image compression," in *Proc. Int. Conf. on Machine Learning (ICML)*, vol. 70, Aug. 2017, pp. 2922–2930.
- [144] D. Minnen, J. Ballé, and G. D. Toderici, "Joint autoregressive and hierarchical priors for learned image compression," in *Advances in Neural Information Processing Systems*, 2018, pp. 10771–10780.
- [145] D. Minnen and S. Singh, "Channel-wise autoregressive entropy models for learned image compression," in *Int. Conf. on Image Compression (ICIP)*, 2020, 2020.
- [146] E. Agustsson, D. Minnen, N. Johnston, J. Balle, S. J. Hwang, and G. Toderici, "Scale-space flow for end-to-end optimized video compression," in *2020 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2020, pp. 8500–8509. [Online]. Available: <https://ieeexplore.ieee.org/document/9157366/>
- [147] C.-Y. Wu, N. Singhal, and P. Krahenbuhl, "Video Compression through Image Interpolation," in *2018 European Conference on Computer Vision (ECCV)*, Sep. 2018, pp. 416–431. [Online]. Available: https://openaccess.thecvf.com/content_ECCV_2018/html/Chao-Yuan_Wu_Video_Compression_through_ECCV_2018_paper.html
- [148] G. Lu, W. Ouyang, D. Xu, X. Zhang, C. Cai, and Z. Gao, "DVC: an end-to-end deep video compression framework," in *2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Sep. 2019, pp. 11 006–11 015.
- [149] O. Rippel, S. Nair, C. Lew, S. Branson, A. G. Anderson, and L. Bourdev, "Learned Video Compression," pp. 3454–3463, Oct. 2019.

- [150] A. Djelouah, J. Campos, S. Schaub-Meyer, and C. Schroers, "Neural inter-frame compression for video coding," in *2019 International Conference on Computer Vision (ICCV)*, Oct. 2019, pp. 6420–6428.
- [151] A. Habibiyan, T. v. Rozendaal, J. M. Tomczak, and T. S. Cohen, "Video compression with rate-distortion autoencoders," in *2019 International Conference on Computer Vision (ICCV)*, Oct. 2019, pp. 7033–7042.
- [152] J. Han, S. Lombardo, C. Schroers, and S. Mandt, "Deep generative video compression," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, Nov. 2019.
- [153] Z. Hu, G. Lu, and D. Xu, "FVC: A new framework towards deep video compression in feature space," in *2021 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 1502–1511.
- [154] B. Liu, Y. Chen, S. Liu, and H.-S. Kim, "Deep learning in latent space for video prediction and compression," in *2021 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 701–710.
- [155] S. Kankanahalli, "End-to-end optimized speech coding with deep neural networks," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 2521–2525.
- [156] W. B. Kleijn, F. S. C. Lim, A. Luebs, J. Skoglund, F. Stimberg, Q. Wang, and T. C. Walters, "Wavenet based low rate speech coding," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE Press, 2018, p. 6766680.
- [157] D. Petermann, S. Beack, and M. Kim, "Harp-net: Hyper-autoencoded reconstruction propagation for scalable neural audio coding," *2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pp. 316–320, 2021.
- [158] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, "Soundstream: An end-to-end neural audio codec," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 495–507, 2022.
- [159] M. C. Anderson, M. S. Atkins, and J. Vaisey, "Task-oriented lossy compression of magnetic resonance images," in *Medical Imaging 1996: Image Display*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Y. Kim, Ed., vol. 2707, Apr. 1996, pp. 239–250.
- [160] L. Pu, M. W. Marcellin, A. Bilgin, and A. Ashok, "Image compression based on task-specific information," in *2014 IEEE International Conference on Image Processing (ICIP)*, 2014, pp. 4817–4821.
- [161] C. Bulla, C. Feldmann, and M. Schink, "Region of interest encoding in video conference systems," in *MMEDIA 2013*, 2013.
- [162] M. Xu, X. Deng, S. Li, and Z. Wang, "Region-of-interest based conversational hevcc coding with hierarchical perception model of face," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 3, pp. 475–489, 2014.
- [163] F. A. Andalo, O. A. B. Penatti, and V. Testoni, "Transmitting what matters: Task-oriented video composition and compression," in *Conference on Graphics, Patterns and Images (SIBGRAPI)*, 2016, pp. 72–79.
- [164] M. Minervini and S. A. Tsafaris, "Classification-aware distortion metric for hevcc intra coding," in *2015 Visual Communications and Image Processing (VCIP)*, 2015, pp. 1–4.
- [165] R. Torfason, F. Mentzer, E. Ágústsson, M. Tschannen, R. Timofte, and L. V. Gool, "Towards image understanding from deep compression without decoding," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=HkXWCMbRW>
- [166] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03385>
- [167] N. Shlezinger and Y. C. Eldar, "Deep task-based quantization," *Entropy*, vol. 23, 2021.
- [168] M. Shohat, G. Tsintsadze, N. Shlezinger, and Y. C. Eldar, "Deep quantization for mimo channel estimation," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 3912–3916.
- [169] M. Jankowski, D. Gunduz, and K. Mikolajczyk, "Wireless image retrieval at the edge," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 1, pp. 89–100, 2021.
- [170] T. Goblick, "Theoretical limitations on the transmission of data from analog sources," *IEEE Transactions on Information Theory*, vol. 11, no. 4, pp. 558–567, 1965.
- [171] M. Gastpar, B. Rimoldi, and M. Vetterli, "To code, or not to code: lossy source-channel communication revisited," *IEEE Transactions on Information Theory*, vol. 49, no. 5, pp. 1147–1158, 2003.
- [172] T. Cover, A. Gamal, and M. Salehi, "Multiple access channels with arbitrarily correlated sources," *IEEE Transactions on Information Theory*, vol. 26, no. 6, pp. 648–657, 1980.
- [173] C. E. Shannon, "Two-way communication channels," in *Proc. 4th Berkeley Symposium Math. Stat. and Prob.*, vol. 1, 1961, pp. 611–644.
- [174] G. Dueck, "A note on the multiple access channel with correlated sources (corresp.)," *IEEE Transactions on Information Theory*, vol. 27, no. 2, pp. 232–235, 1981.
- [175] W. Kang and S. Ulukus, "An outer bound for multiple access channels with correlated sources," in *2006 40th Annual Conference on Information Sciences and Systems*, 2006, pp. 240–244.
- [176] J. Nayak, E. Tuncel, and D. Gunduz, "Wyneráziv coding over broadcast channels: Digital schemes," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1782–1799, 2010.
- [177] S. Bross, A. Lapidoth, and S. Tinguely, "Broadcasting correlated Gaussians," *IEEE Trans. Inf. Theory.*, vol. 56, no. 7, pp. 3057–3068, Jul. 2010.
- [178] A. Lapidoth and S. Tinguely, "Sending a bivariate Gaussian over a Gaussian MAC," *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2714–2752, Jun. 2010.
- [179] —, "Sending a bivariate gaussian source over a gaussian mac with feedback," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1852–1864, Apr. 2010.
- [180] I. Estella-Agueri and D. G. Á. and Á. V. Z., "Gaussian joint source-channel coding for the strong interference channel," in *2015 IEEE Information Theory Workshop (ITW)*, 2015, pp. 1–5.
- [181] I. Estella Aguerri and D. Gunduz, "Correlated Gaussian sources over Gaussian weak interference channels," in *2015 IEEE Information Theory Workshop - Fall (ITW)*, 2015, pp. 84–88.
- [182] I. Estella Aguerri and D. G. Á. and Á. V. Z., "Joint source-channel coding with time-varying channel and side-information," *IEEE Transactions on Information Theory*, vol. 62, no. 2, pp. 736–753, 2016.
- [183] S. Sreekumar and D. Gunduz, "Distributed hypothesis testing over discrete memoryless channels," *IEEE Transactions on Information Theory*, vol. 66, no. 4, pp. 2044–2066, 2020.
- [184] H. Shimokawa, T. S. Han, and S. Amari, "Error bound of hypothesis testing with data compression," in *Proc. IEEE Int. Symp. Inf. Theory, Trondheim, Norway*, 1994.
- [185] S. Borade, B. Nakiboğlu, and L. Zheng, "Unequal error protection: An information-theoretic perspective," *IEEE Trans. on Info. Theory*, vol. 55, no. 12, pp. 5511–5539, Dec. 2009.
- [186] P. Minero, S. H. Lim, and Y.-H. Kim, "A unified approach to hybrid coding," *IEEE Transactions on Information Theory*, vol. 61, no. 4, pp. 1509–1523, 2015.
- [187] J.-F. Chamberland and V. Veeravalli, "Asymptotic results for decentralized detection in power constrained wireless sensor networks," *IEEE Journal on Selected Areas in Communications*, vol. 22, no. 6, pp. 1007–1015, 2004.
- [188] R. Niu, B. Chen, and P. Varshney, "Fusion of decisions transmitted over rayleigh fading channels in wireless sensor networks," *IEEE Transactions on Signal Processing*, vol. 54, no. 3, pp. 1018–1027, 2006.
- [189] T. Duman and M. Salehi, "Decentralized detection over multiple-access channels," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, no. 2, pp. 469–476, 1998.
- [190] G. Mergen and L. Tong, "Type based estimation over multiaccess channels," *IEEE Transactions on Signal Processing*, vol. 54, no. 2, pp. 613–626, 2006.
- [191] C.-Z. Lee, L. P. Barnes, and A. x00D6; zgx00FC; r, "Over-the-air statistical estimation," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 2, pp. 548–561, 2022.
- [192] G. Cheung and A. Zakhor, "Bit allocation for joint source/channel coding of scalable video," *IEEE Trans. on Image Processing*, vol. 9, no. 3, pp. 340–356, Mar. 2000.
- [193] N. Farvardin, "A study of vector quantization for noisy channels," *IEEE Transactions on Information Theory*, vol. 36, no. 4, pp. 799–809, July 1990.
- [194] N. Farvardin and V. Vaishampayan, "On the performance and complexity of channel-optimized vector quantizers," *IEEE Transactions on Information Theory*, vol. 37, no. 1, pp. 155–160, Jan 1991.
- [195] M. Skoglund, "Soft decoding for vector quantization over noisy channels with memory," *IEEE Trans. on Information Theory*, vol. 45, no. 4, pp. 1293–1307, May 1999.
- [196] I. Kozintsev and K. Ramchandran, "Robust image transmission over energy-constrained time-varying channels using multiresolution joint source-channel coding," *IEEE Trans. on Signal Processing*, vol. 46, no. 4, pp. 1012–1026, Apr. 1998.
- [197] M. Skoglund, "On channel-constrained vector quantization and index assignment for discrete memoryless channels," *IEEE Transactions on Information Theory*, vol. 45, no. 7, pp. 2615–2622, Nov 1999.

- [198] S. Jakubczak and D. Katabi, "Softcast: One-size-fits-all wireless video," *SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 4, p. 449&450, aug 2010.
- [199] O. Y. Bursalioğlu, G. Caire, and D. Divsalar, "Joint source-channel coding for deep space image transmission using rateless codes," in *2011 Information Theory and Applications Workshop*, 2011, pp. 1–10.
- [200] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," California Univ San Diego La Jolla Inst for Cognitive Science, Tech. Rep., 1985.
- [201] E. Boursoulazte, D. Burth Kurka, and D. Gunduz, "Deep joint source-channel coding for wireless image transmission," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 3, pp. 567–579, Sep. 2019.
- [202] N. Farsad, M. Rao, and A. Goldsmith, "Deep learning for joint source-channel coding of text," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 2326–2330.
- [203] D. B. Kurka and D. Gunduz, "Deepjssc-f: Deep joint source-channel coding of images with feedback," *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 178–193, 2020.
- [204] D. B. Kurka and D. Gunduz, "Bandwidth-agile image transmission with deep joint source-channel coding," *IEEE Transactions on Wireless Communications*, vol. 20, no. 12, pp. 8081–8095, 2021.
- [205] T.-Y. Tung and D. Gündüz, "DeepWiVe: Deep-learning-aided wireless video transmission," *ArXiv*, vol. abs/2111.13034, 2021.
- [206] B. Guler and A. Yener, "Semantic index assignment," in *2014 IEEE International Conference on Pervasive Computing and Communication Workshops (PERCOM WORKSHOPS)*, 2014, pp. 431–436.
- [207] B. Guler, A. Yener, and A. Swami, "The semantic communication game," *IEEE Transactions on Cognitive Communications and Networking*, vol. 4, no. 4, pp. 787–802, 2018.
- [208] Z. Weng and Z. Qin, "Semantic communication systems for speech transmission," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 8, pp. 2434–2444, 2021.
- [209] H. Xie, Z. Qin, X. Tao, and K. B. Letaief, "Task-oriented multi-user semantic communications," *arXiv preprint arXiv:2112.10255*, Dec. 2021.
- [210] Z. Weng, Z. Qin, and G. Y. Li, "Semantic communications for speech recognition," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Madrid, Spain, Dec. 2021.
- [211] M. Schuster and K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Trans. Signal Process.*, vol. 45, no. 11, pp. 2673–2681, Nov. 1997.
- [212] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proc. Int. Conf. Mach. Learning (ICML)*, Pittsburgh, USA, Jun. 2006, pp. 369–376.
- [213] T. Han, Q. Yang, Z. Shi, S. He, and Z. Zhang, "Semantic-aware speech to text transmission with redundancy removal," *arXiv preprint arXiv:2202.03211*, Feb. 2022.
- [214] H. Seo, J. Park, M. Bennis, and M. Debbah, "Semantics-native communication with contextual reasoning," *arXiv preprint arXiv:2108.05681*, Aug. 2021.
- [215] M. K. Farshbafan, W. Saad, and M. Debbah, "Common language for goal-oriented semantic communications: A curriculum learning framework," *arXiv preprint arXiv:2111.08051*, Nov. 2021.
- [216] F. Liu, W. Tong, Z. Sun, and C. Guo, "Task-oriented semantic communication systems based on extended rate-distortion theory," *arXiv preprint arXiv:2201.10929*, Jan. 2022.
- [217] J. Shao, Y. Mao, and J. Zhang, "Learning task-oriented communication for edge inference: An information bottleneck approach," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 1, pp. 197–211, Nov. 2022.
- [218] D. Gunduz, D. B. Kurka, M. Jankowski, M. M. Amiri, E. Ozfatura, and S. Sreekumar, "Communicate to learn at the edge," *IEEE Communications Magazine*, vol. 58, no. 12, pp. 14–19, 2020.
- [219] D. G. M. Jankowski and K. Mikołajczyk, "Airmet: Neural network transmission over the air," in *Proc. IEEE Int. Symp. Inf. Theory*, Helsinki, Finland, 2022.
- [220] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. Graf, "Pruning Filters for Efficient Vonnets," in *Int'l Conf. on Learning Repr. (ICLR)*, 2017.
- [221] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," in *International Conference on Learning Representations (ICLR)*, 2019.
- [222] V. A. Kotelnikov, *The Theory of Optimum Noise Immunity*. McGraw-Hill Book Company, Inc., 1959.
- [223] Y. Shao, S. C. Liew, and D. Gündüz, "Denosing noisy neural networks: A bayesian approach with compensation," *arXiv:2105.10699*, 2021.
- [224] M. M. Amiri and D. Gunduz, "Machine learning at the wireless edge: Distributed stochastic gradient descent over-the-air," *IEEE Transactions on Signal Processing*, vol. 68, pp. 2155–2169, Mar. 2020.
- [225] —, "Federated learning over wireless fading channels," *IEEE Transactions on Wireless Communications*, vol. 19, no. 5, pp. 3546–3557, May 2020.
- [226] G. Zhu, Y. Du, D. Gunduz, and K. Huang, "One-bit over-the-air aggregation for communication-efficient federated edge learning: Design and convergence analysis," *IEEE Transactions on Wireless Communications*, pp. 1–1, 2020.
- [227] M. M. Amiri, T. M. Duman, D. Gündüz, S. R. Kulkarni, and H. V. Poor, "Blind federated edge learning," *IEEE Transactions on Wireless Communications*, vol. 20, no. 8, pp. 5129–5143, 2021.
- [228] E. Ozfatura, S. Rini, and D. Gündüz, "Decentralized sgd with over-the-air computation," in *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, 2020, pp. 1–6.
- [229] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society, Series B*, vol. 39, no. 1, pp. 1–38, 1977.
- [230] R. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Transactions on Information Theory*, vol. 18, no. 4, pp. 460–473, Jul 1972.
- [231] S. Arimoto, "An algorithm for computing the capacity of arbitrary discrete memoryless channels," *IEEE Trans. Inf. Theory*, vol. IT-18, no. 1, pp. 12 – 20, Jan. 1972.
- [232] A. Globerson and N. Tishby, "On the optimality of the Gaussian information bottleneck curve," *technical report, Hebrew University*, 2004.
- [233] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss, "Information bottleneck for Gaussian variables," *Journal of Machine Learning Research*, vol. 6, pp. 165–188, Feb. 2005.
- [234] D. P. Kingma and M. Welling, "Auto-encoding variational bayes." *CoRR*, vol. abs/1312.6114, 2013. [Online]. Available: <http://dblp.uni-trier.de/db/journals/corr/corr1312.html#KingmaW13>
- [235] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "β-vae: Learning basic visual concepts with a constrained variational framework," 2016.
- [236] 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, 2021.
- [237] F. Pezone, S. Barbarossa, and P. Di Lorenzo, "Goal-oriented communication for edge learning based on the information bottleneck," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 8832–8836.
- [238] C. J. Maddison, A. Mnih, and Y. W. Teh, "The concrete distribution: A continuous relaxation of discrete random variables," *CoRR*, vol. abs/1611.00712, 2016. [Online]. Available: <http://arxiv.org/abs/1611.00712>
- [239] E. Jiang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," 2017. [Online]. Available: <https://arxiv.org/abs/1611.01144>
- [240] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization." *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://dblp.uni-trier.de/db/journals/corr/corr1412.html#KingmaB14>
- [241] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *EEE Journal on Selected Areas in Commun.*, vol. 39, no. 5, pp. 1183–1210, May 2021.
- [242] A. Kosta, N. Pappas, and V. Angelakis, "Age of information: A new concept, metric, and tool," *Foundations and Trends in Networking*, vol. 12, no. 3, pp. 162–259, 2017.
- [243] M. A. Abd-Elmagid, N. Pappas, and H. S. Dhillon, "On the role of age of information in the Internet of things," *IEEE Commun. Magazine*, vol. 57, no. 12, pp. 72–77, Dec. 2019.
- [244] Y. Sun, I. Kadota, R. Talak, and E. Modiano, "Age of information: A new metric for information freshness," *Synthesis Lectures on Communication Networks*, vol. 12, no. 2, pp. 1–224, 2019.
- [245] R. D. Yates and S. K. Kaul, "The age of information: Real-time status updating by multiple sources," *IEEE Trans. on Info. Theory*, vol. 65, no. 3, pp. 1807–1827, Mar. 2019.
- [246] R. D. Yates, "The age of information in networks: Moments, distributions, and sampling," *IEEE Trans. on Info. Theory*, vol. 66, no. 9, pp. 5712–5728, Sept. 2020.
- [247] A. Soysal and S. Ulukus, "Age of information in G/G/1/1 systems: Age expressions, bounds, special cases, and optimization," *IEEE Trans. on Info. Theory*, to appear.

- [248] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Trans. on Info. Theory*, vol. 62, no. 4, pp. 1897–1910, Apr. 2016.
- [249] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "On the age of information with packet deadlines," *IEEE Trans. on Info. Theory*, vol. 64, no. 9, pp. 6419–6428, Sept. 2018.
- [250] M. Moltafet, M. Leinonen, and M. Codreanu, "On the age of information in multi-source queueing models," *IEEE Trans. on Commun.*, vol. 68, no. 8, pp. 5003–5017, Aug. 2020.
- [251] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "Age of information performance of multiaccess strategies with packet management," *Journal of Commun. and Networks*, vol. 21, no. 3, pp. 244–255, June 2019.
- [252] Y. Inoue, H. Masuyama, T. Takine, and T. Tanaka, "A general formula for the stationary distribution of the age of information and its application to single-server queues," *IEEE Trans. Info. Theory*, vol. 65, no. 12, pp. 8305–8324, Dec. 2019.
- [253] J. P. Champati, H. Al-Zubaidy, and J. Gross, "On the distribution of AoI for the GI/GI/1/1 and GI/GI/1/2* systems: Exact expressions and bounds," in *Proc., IEEE INFOCOM*, 2019.
- [254] M. A. Abd-Elmagid and H. S. Dhillon, "Closed-form characterization of the MGF of AoI in energy harvesting status update systems," *IEEE Trans. on Info. Theory*, to appear.
- [255] —, "Age of information in multi-source updating systems powered by energy harvesting," *IEEE Journal on Selected Areas in Info. Theory*, vol. 3, no. 1, pp. 98–112, Mar. 2022.
- [256] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Transactions on Networking*, vol. 28, no. 5, pp. 2215–2228, 2020.
- [257] T. Linder and S. Yüksel, "On optimal zero-delay coding of vector markov sources," *IEEE Trans. on Info. Theory*, vol. 60, no. 10, pp. 5975–5991, 2014.
- [258] A. Khina, V. Kostina, A. Khisti, and B. Hassibi, "Sequential coding of gauss-markov sources with packet erasures and feedback," in *IEEE Info. Theory Workshop (ITW)*. IEEE, 2017.
- [259] O. Orhan, D. Gündüz, and E. Erkip, "Source-channel coding under energy, delay, and buffer constraints," *IEEE Trans. on Wireless Commun.*, vol. 14, no. 7, pp. 3836–3849, 2015.
- [260] 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.
- [261] T.-Y. Tung, S. Kobus, and D. Gündüz, "Context-aware effective communications," in *2021 55th Asilomar Conference on Signals, Systems, and Computers*, 2021, pp. 334–339.
- [262] A. Mostaani, O. Simeone, S. Chatzinotas, and B. Ottersten, "Learning-based Physical Layer Communications for Multiagent Collaboration," in *IEEE Int'l Symp. on Personal, Indoor and Mobile Radio Comms. (PIMRC)*, Sep. 2019, pp. 1–6.
- [263] J. N. Foerster, Y. M. Assael, N. de Freitas, and S. Whiteson, "Learning to Communicate with Deep Multi-Agent Reinforcement Learning," *arXiv:1605.06676 [cs]*, May 2016, arXiv: 1605.06676.
- [264] J. Jiang and Z. Lu, "Learning attentional communication for multi-agent cooperation," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, p. 7265–7275.
- [265] N. Jaques, A. Lazaridou, E. Hughes, C. Gulcehre, P. Ortega, D. Strouse, J. Leibo, and N. deFreitas, "Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning," *arXiv:1810.08647 [cs, stat]*, Jun. 2019. [Online]. Available: <http://arxiv.org/abs/1810.08647>
- [266] A. Das, T. Gervet, J. Romoff, D. Batra, D. Parikh, M. Rabbat, and J. Pineau, "TarMAC: Targeted multi-agent communication," in *36th Int'l Conf. on Machine Learning*. Long Beach, California: PMLR, Jun 2019, pp. 1538–1546.