

Causal Semantic Communication for Digital Twins: A Generalizable Imitation Learning Approach

Christo Kurisummoottil Thomas, *Member, IEEE* and Walid Saad, *Fellow, IEEE*,
and Yong Xiao, *Senior Member, IEEE*

Abstract—A digital twin (DT) leverages a virtual representation of the physical world, along with communication (e.g., 6G), computing (e.g., edge computing), and artificial intelligence (AI) technologies to enable many connected intelligence services. In order to handle the large amounts of network data based on digital twins (DTs), wireless systems can exploit the paradigm of semantic communication (SC) for facilitating informed decision-making under strict communication constraints by utilizing AI techniques such as causal reasoning. In this paper, a novel framework called causal semantic communication (CSC) is proposed for DT-based wireless systems. The CSC system is posed as an imitation learning (IL) problem, where the transmitter, with access to optimal network control policies using a DT, teaches the receiver using SC over a bandwidth-limited wireless channel how to improve its knowledge to perform optimal control actions. The causal structure in the transmitter’s data is extracted using novel approaches from the framework of deep end-to-end causal inference, thereby enabling the creation of a semantic representation that is causally invariant, which in turn helps generalize the learned knowledge of the system to new and unseen situations. The CSC decoder at the receiver is designed to extract and estimate semantic information while ensuring high semantic reliability. The receiver control policies, semantic decoder, and causal inference are formulated as a bi-level optimization problem within a variational inference framework. This problem is solved using a novel concept called network state models, inspired from world models in generative AI, that faithfully represents the environment dynamics leading to data generation. Furthermore, the proposed framework includes an analytical characterization of the performance gap that results from employing a suboptimal policy learned by the receiver that uses the transmitted semantic information to construct a model of the physical environment. The CSC system utilizes two concepts, namely the integrated information theory principle in the theory of consciousness and the abstract cell complex concept in topology, to precisely express the information content conveyed by the causal states and their relationships. Through this analysis, novel formulations of semantic information, semantic reliability, distortion, and similarity metrics are proposed,

Christo Kurisummoottil Thomas and Walid Saad are with the Wireless@VT, Bradley Department of Electrical and Computer Engineering, Virginia Tech, Arlington, VA, USA.(emails:{christokt,walids}@vt.edu).

Yong Xiao is with the School of Electronic Information and Communications at the Huazhong University of Science and Technology, Wuhan 430074, China, also with the Peng Cheng Laboratory, Shenzhen, Guangdong 518055, China, and also with the Pazhou Laboratory (Huangpu), Guangzhou, Guangdong 510555, China (e-mail: yongxiao@hust.edu.cn).

The work of Christo Kurisummoottil Thomas and Walid Saad are supported by US National Science Foundation under Grant CNS-2210254.

which extend beyond Shannon’s concept of uncertainty. Simulation results demonstrate that the proposed CSC system outperforms conventional wireless and state-of-the-art SC systems by achieving better semantic reliability with reduced bits and enabling better control policies over time thanks to the generative AI architecture.

I. INTRODUCTION

Digital twins (DTs) are replicas of the physical world [1] created through the use of simulation software, data analytics, and sensor data. These virtual replicas rely on real-time data and advanced algorithms to model and predict the behavior and performance of physical systems, enabling better decision-making and improved efficiency across a range of industries, such as manufacturing, healthcare, transportation, and aerospace. DTs can also be used to assist wireless systems [2] to enable self-configuration, proactive online learning, and support for *connected intelligence (CI)* applications such as haptics, brain-computer interaction, flying vehicles, extended reality (XR), and the metaverse. In order to create a DT within a wireless system (e.g., at a base station or mobile edge computing server), there is need for significant computing resources and transmission of large volumes of data. Additionally, the DT data collected at the base station will be employed to develop a model of the physical world, which can assist in making better control decisions or aiding CI applications. Transmitting such a staggering amount of information back from the base station to the end users can be highly inefficient from a resource utilization perspective, leading to increased latency, high power consumption, and reduced spectrum efficiency. To meet the demands of high-rate, high-reliability, and time-criticality for the aforementioned CI applications, DT-powered 6G systems could transmit only essential information relevant to the end-user. This concept forms the basis of *semantic communication (SC)* systems [3]–[5]. Clearly, integrating DTs with SC offers a more effective solution for handling large amounts of data and optimizing resource utilization. Moreover, by employing SC, DTs can leverage contextual knowledge to better interpret and process data generated from diverse sources, leading to more precise predictions for event monitoring, optimal control decisions, and improved au-

tonomous agent capabilities. However, integrating DTs with SCs poses unique *challenges* particularly on the SC side. For instance, the transmitter that manages the DT must ensure that the model is trained on representative data, which can be challenging in non-stationary and dynamic wireless environments. In addition, transmitting data from a DT-enabled transmitter to facilitate real-time prediction, control, or reconstruction tasks at the end-users can be challenging due to the need for ultra-low-latency and ultra-reliable communication. To address the challenges of DT-based semantic communication, one promising approach is to extract the causal structure inherent in the network data. This enables the development of an accurate physical environment model at the transmit and receive nodes with fewer training samples. By leveraging generative artificial intelligence (AI) methods, this approach ultimately enables real-time prediction, reconstruction, and control.

A. Prior Works

Despite recent AI-based SC designs [6]–[11], these prior works failed to address a few critical aspects that hinder the integration of DT and SC. Firstly, most of these designs ignore the need for a rigorous formulation of semantics, instead considering the semantic concepts as being mapped to data via a probabilistic transformation [6]. For a DT, this probabilistic design could lead to inaccurate physical models at the end nodes, which in turn would lead to inaccurate decisions. Second, the solutions of [6]–[11] suffer from limited generalizability to scenarios not encountered during training. These approaches typically use variational auto-encoders (VAEs) or transformers without structure extraction, which may reduce the ability to handle unseen wireless environments. This lack of generalizability necessitates significant retraining efforts, leading to increased communication overheads and delays, which can hinder real-time prediction, reconstruction, and control under communication constraints for DT-based wireless systems. Meanwhile, the works in [12]–[17] investigated the use of DTs to enhance the performance of wireless systems. For instance, the authors in [12] proposed a distributed DT-powered federated learning framework to support edge computing in industrial IoT, while the work in [13] suggested constructing DTs at edge networks by using blockchain and federated learning. In a recent work [14], the authors investigated the joint synchronization of DTs and sub-metaverses within a distributed metaverse framework. Despite being interesting, the prior art [12]–[17] has not adequately addressed modeling the environment dynamics that capture the data observed at the DT side. The only exception is the work in [18] that proposed the use of a DT that builds a Bayesian model of the communication system environment using conventional signal processing techniques. While the Bayesian modeling approach presented is intriguing, the solution of [18] is not generalizable, i.e., it is limited to the specific data distribution used to create the DT, and thus it cannot generalize

to multiple wireless environment. Moreover, an important aspect missing in the prior art [12]–[18] is the efficient transmission of network state and control information from DT-based nodes to edge users, which poses a challenge to achieving the high-reliability and low-latency goals of future wireless networks. To the authors' knowledge, this work will be the *first* to combine SC and DTs to address the aforementioned challenges.

B. Contributions

In contrast to the state-of-art that lacks a rigorous definition of semantics and generalizable SC system design, the main contribution of this paper is a novel framework for designing a causal SC (CSC) system consisting of two components: 1) a transmitter (called teacher) based on DTs that identifies the semantic content elements (SCEs) present in the data and performs causal discovery of the state transitions being transmitted, and 2) a receiver node (acting as apprentice) that learns the environment dynamics using the history of transmitted semantic representations and designs better control policies to maximize semantic effectiveness. While imitation learning (IL) [19] is a promising method for implementing this envisioned teacher-apprentice framework for SC [20], practical challenges arise when implementing existing IL methods [21] in DT-based SC systems. These challenges stem from the resource-intensive nature of demonstration data available at the DT side and the imperfections that may arise when transmitting data over a wireless channel, leading to inaccurate policy learning. The authors in [20] used model-free reinforcement learning (MFRL) to infer implicit semantic entities and relations from explicit semantics or observables. However, the testing performance of this approach is dependent on the specific data distribution used during training, and thus not generalizable. To overcome these issues, we propose to advance IL utilizing model-based RL (MBRL) thereby enabling the development of a model of the wireless environment dynamics using structural causal models (SCM) [22]. The proposed AI-based components in the SC chain rely on novel semantic information measures inspired by the concept of *integrated information theory (IIT)* in the literature on the *theory of consciousness* [23]. IIT measures help identify distinct *SCEs*, which hold meanings or semantics and are present within the network states observed by the DT. In summary, our key contributions include:

- We introduce *novel information measures* for the learned *SCM* at the imitator, inspired from the IIT measures [24]. We also propose a new semantic state abstraction concept that utilizes the intrinsic information concept from IIT. Semantic state abstraction plays a crucial role as it helps eliminate irrelevant information observed at the transmitter side, thus reducing transmitted information. Moreover, it enhances the ability of the agent to generalize to previously unseen areas of the state space, which could correspond to

different wireless environments. This allows us to develop generalizable native AI-based wireless SC systems.

- We use the concept of abstract simplicial complexes in topology to theoretically characterize causal relations among SCEs in the data observed at the DT. Moreover, we demonstrate that the SCEs and their associations can be algebraically defined as a cell complex structure. By defining the SCEs in accordance with IIT, these topological characterizations enable us to define semantic metrics such as similarity and information that go beyond conventional information theory concepts. This capability enables the creation of DT-based SC systems that can precisely identify the causal structure and consequently develop causally invariant semantic representations.
- At the receiver, we create a “*network state model*” similar to the concept of world models in generative AI [25]. We design a semantic decoder that extracts the maximum semantic information. However, due to limited computing resources at the imitator node, the model learns a sub-optimal version of the environment dynamics. This model is then used to learn imitator policies for controlling communication tasks. To solve for the neural network (NN) parameters, we propose a bi-level optimization method within an MBRL framework. This approach enables us to learn both the state transition model and the imitator policy.
- We analytically characterize the performance shortfall (compared to the expert agent’s policy) in terms of the quality-of-experience (QoE) associated with the sub-optimal network state models created at the receiver node. The QoE is defined as a function of the semantic effectiveness resulting from the transmitted semantic representation.
- Simulation results demonstrate the superiority of the proposed CSC in semantic reliability and throughput delivered by a DT-based SC system, in contrast to similar wireless systems that rely on traditional AI approaches. Furthermore, even when the channel quality is very poor, the proposed CSC achieves a semantic reliability that is 3.4 times better than classical DT-based SC systems due to its generative AI architecture at the receiver that allows it to reconstruct data. The simulations conducted in our study provide further evidence that the proposed CSC can adapt quickly to non-stationary conditions in wireless environments. The CSC requires approximately 3K fewer samples for retraining, which validates its generalizability.

The rest of this paper is organized as follows. In Section II, we present the proposed CSC system model. Section III describes the proposed IIT measures for quantifying the semantic information. Section IV introduces the MBRL based IL for learning the transmitter and receiver modules. Section IV provides simulation results. Finally, conclusions

are drawn in Section VI¹.

Notations: Lower-case letter a is a scalar, boldface lower-case \mathbf{a} and upper-case letter \mathbf{A} represent a vector and a matrix, respectively. A set (either discrete or continuous entries) or a topological space is represented using Calligraphic font \mathcal{X} . \mathcal{R}^M represents an M -dimensional vector whose entries belong to real numbers \mathcal{R} . $[a, b]$ represents the real number range between a and b . $\mathbf{A} \odot \mathbf{A}$ represents the element-wise multiplication of two matrices. $|a|$ represents the absolute value of a .

II. SYSTEM MODEL

Consider a point-to-point SC system that employs a DT to aid in making control decisions. This system can capture many practical applications. For example, in the context of 6G and open radio access network (ORAN) systems, a network DT located at a cloud data center would replicate various latent aspects of a network, including signals, coverage, interference, traffic behavior, and user mobility, across different frequency layers. The DT optimizes sensitive parameters, such as radiated power or multi-user scheduling, by providing a safe simulation environment without any actual risks to the real network. Within the ORAN architecture, the radio unit (RU) component is located at the base station and has limited computational power [26]. To compensate for this, the RU relies on the transmission of information from the control unit (CU) component that implements the DT. The CU has more knowledge about the network architecture and user profiles, aided by the cloud data center, and can transmit this information to the RU via a fronthaul wireless link. By implementing a network DT-based SC system over the fronthaul link, the network’s efficiency (bandwidth utilization) and security can be significantly enhanced, making it an essential tool for network management. Another key example that is captured by our model is in CI applications, where DTs can improve user experience in various autonomous wireless services, such as robotic surgery, telehealth, driverless vehicles, and industrial robotic manufacturing sites. By creating a virtual representation of the physical system, a DT can simulate and predict its behavior in real-time, enabling proactive decision-making and optimized system performance to enhance the user experience. For instance, a DT can reduce latency and ensure reliable communication in telehealth services, or enhance safety and efficiency in driverless vehicles and industrial robotic manufacturing sites.

The functionalities of the DT are summarized in Fig. 1. The DT is created by an expert agent who analyzes observations from various sensing elements (defined as (\mathbf{s}^t, θ)) in the wireless environment. These observations are used

¹This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org/>, provided by the author. The material includes all the appendices that derive the proofs of Lemmas and Theorems in the paper. Contact christokt@vt.edu for further questions about this work.

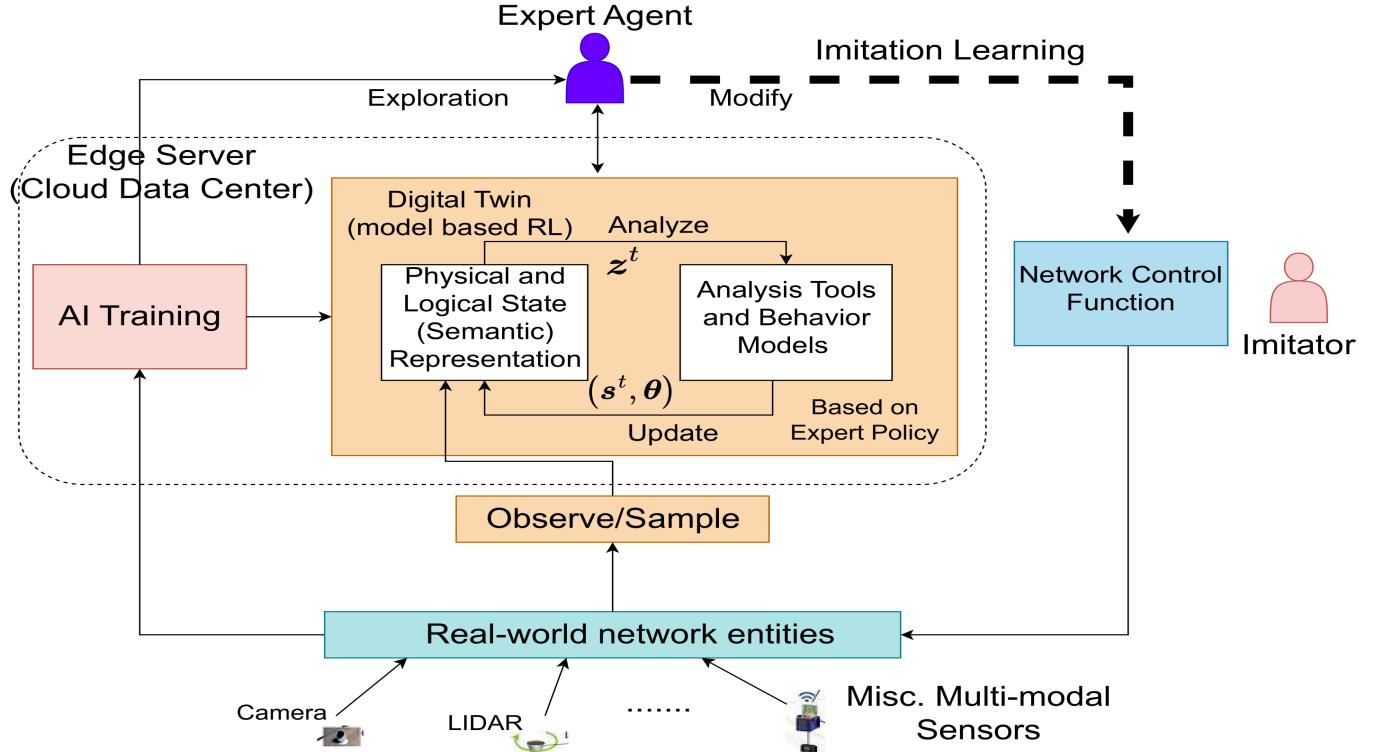


Fig. 1: Illustration of our digital twin model showing expert (transmitter that has access to different sensor observations) and imitator agent interactions.

to model the physical and logical state of the network represented by the transition probability $p(s^{t+1} | s^t, \theta, a^t)$, which is then used to formulate optimal control decisions a^t . The expert agent communicates either a condensed version of the network state or the optimal control decisions defined as z^t to the imitator agent. Subsequently, the imitator extracts the state information and performs the necessary low-level actions. The interactions between expert and imitator agents in the SC system, as described here, can be classified as an example of IL. By utilizing a DT, accurate physical system models and real-time data can be made available to inform decision-making, facilitating effective control decisions. We next explain the details of our DT-based SC system as shown in Fig. 2. The source side of the transmission consists of a resource-constrained node aided by an edge server with significant computational and memory capabilities. The source user and the edge server together act as the expert agent (that implements the DT) to teach an apprentice or imitator agent, serving as the receiver, with no prior knowledge of the expert's data. Here, one naturally wonders whether communicating only the expert actions determined using the policy evaluated by the DT could be sufficient. While communicating only actions may be relevant when the imitator agent has limited capabilities, for cases in which the agent must possess more advanced intelligence, additional information beyond actions may be necessary for effective imitation. In the case of a SC system, building an intelligent imitator agent is

crucial. For instance, during channel outages lasting over several communication time slots, the receiver must still generate policy information by leveraging the historical states received over the network and potentially using limited sensing information. Merely communicating the actions does not offer significant information about the observed states at the transmitter, which, in turn, fails to facilitate the environment modeling at the imitator. This, in turn, motivates designing systems that communicate meaningful semantic representations of the states observed at the expert agent to the imitator. Further, the imitator can utilize the potential of generative AI algorithms [27] to reconstruct the network states. Next, we outline the different AI-based components present in the proposed SC system.

A. Causal imitation learning for digital twin-based SC systems

Given the IL model, we implement the SC functionalities as follows. We divide the transmissions into training and communication phases. During training, the source user samples random initial states from a particular distribution, $p(s^0)$. The sampler implements this functionality during training while the states are observable from the environment during the communication phase. Given the initial state, the state transition is determined by the action of the source, and it is sampled using the *expert policy*, $a^t \sim \pi_E(a^t | s^t, \theta)$. The optimal action a^t depends on the state at the current time t and *confounding variables*

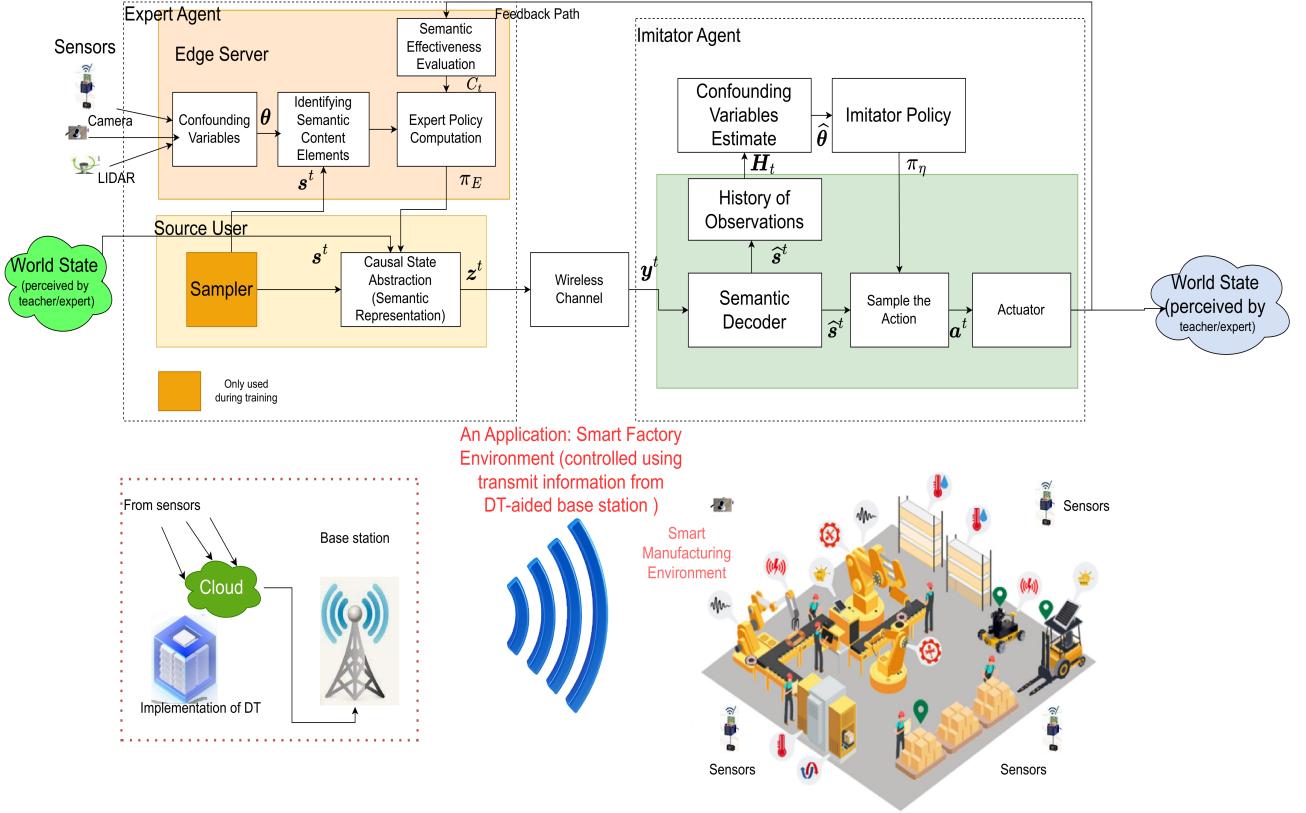


Fig. 2: Proposed SC system framework based on IL. The figure also illustrates an example application as an autonomous wireless system that includes a DT designed explicitly for a smart manufacturing environment. The application is based on sensing information gathered at the central node and is tailored to the proposed system model.

(from various sensing devices) θ that could represent background information about the communication environment accessible only to the edge server. The expert agent computes the future state $s^{t+1} \sim p(s^{t+1} | s^t, a^t, \theta)$, using the information from θ and the current state s_t . In our DT-based system, efficient transmission is of utmost importance due to the vast amounts of data the expert agent processes. To address this issue, the expert agent could exploit a *causal discovery AI* method [28] to extract data structure (inherent in the state transition dynamics) while identifying the SCEs. In this setup, a potential question is how to identify the SCEs (defined as u^t) comprised in any subset of the state s_i^t . SCEs are entities in the source data that hold *meaning (semantics)* or significance, aiming to improve the knowledge learned by the imitator. We aim to identify the SCEs in the data and characterize the semantic information in the causal structure. Causality is apt here since it helps with generalization in SC. It allows the communicating nodes to understand the underlying mechanisms and processes that drive the relationships between network variables rather than just observing their associations. This deeper understanding of the underlying causal relationships can help the nodes to make more accurate predictions or control decisions and develop more robust models that can generalize well to new, unseen data. To represent the causal

structure that underlies the transition from s^t to s^{t+1} , our CSC system employs an MBRL approach. The causal model underlying the state transitions based on Pearl's representation of causality [22] is defined next.

Definition 1. An *SCM* is a collection of elements represented as $< \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) >$, where \mathbf{V} represents endogenous variables (cause and effect variables), and \mathbf{U} represents exogenous variables (random, unknown noise). The set of structural functions \mathcal{F} maps each $f_v \in \mathcal{F}$ to $v \in \mathbf{V}$ such that $f_v(Pa_v, \mathbf{u}_v)$ determines v where the set of parents $Pa_v \subset \mathbf{V}$ and $\mathbf{u}_v \subset \mathbf{U}$. The exogenous distribution $P(\mathbf{U})$ determines the values of \mathbf{U} , and thus the distribution of endogenous variables \mathbf{V} .

The apprentice (imitator) can observe only a subset of endogenous variables, partitioned into \mathbf{O} and \mathbf{L} , where $\mathbf{O} = s^t$ (in reality, an encoded version communicated by the teacher) and $\mathbf{L} = \theta$ (learned by the apprentice). Here, $s^t \in \mathcal{R}^D$. The marginal distribution $P(\mathbf{O})$ is called the observational distribution. MBRL allows communicating the *causal state abstraction* (i.e., the "semantic representation") z^t at any given time over a wireless channel susceptible to errors. The confounding variables are not transmitted due to the communication and memory constraints associated with the wireless channel and imitator, respectively. The

design of the semantic encoder (that maps the states to actual transmit signals) must be done in a way to reduce the amount of data transmitted while maintaining *semantic effectiveness*, defined as a measure of the accuracy of the actions performed using the semantic reconstruction process at the imitator's end.

1) *MBRL Model*: Next, we formulate the problem of expert policy computation as sequential decision-making in a finite-horizon episodic Markov decision process (MDP), defined as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, p(\mathbf{s}^0), H, \boldsymbol{\theta})$. Here, $\boldsymbol{\theta}$ is a vector of confounding variables that the imitator does not have access to and $P(\boldsymbol{\theta})$ is the probability distribution of $\boldsymbol{\theta}$. $\boldsymbol{\theta}$ is fixed within each episode but can vary across different episodes. The M -dimensional vector $\boldsymbol{\theta}$ represents the sensing information available from M sensors which may provide different modality information to supplement the optimal policy. The MDP framework comprises a set of states \mathcal{S} , a set of actions \mathcal{A} , a deterministic reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ that provides evaluative feedback to the imitator agent, a transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ that describes the distribution of future states, an initial state distribution $p(\mathbf{s}^0) \in \Delta(\mathcal{S})$, and a maximum episode length or horizon H . The unit interval constraint for the reward function is adopted for ease of analysis, and our framework can handle general reward functions. The imitator agent cannot access neither the transition function \mathcal{T} nor the reward function \mathcal{R} ; hence, both are treated as random variables (known as the *environment*). The N -dimensional causal state description $\mathbf{s} \in \mathcal{S} \subset \mathcal{R}^N$ is extracted from raw observations and represented as an SCM. Let \mathcal{M}^* be the true MDP that the agent interacts with and attempts to solve over K episodes, with the environment model $\mathcal{E} = (\mathcal{R}^*, \mathcal{T}^*, \boldsymbol{\theta})$. However, the imitator agent's computing capabilities and time constraints limit the accuracy of the true environment that can be learned, resulting in a learned MDP $\widehat{\mathcal{M}}$ (the respective learned environment is $\widehat{\mathcal{E}}$). Each episode involves the agent taking exactly H steps, starting from an initial state $\mathbf{s}^0 \sim p(\mathbf{s}^0)$. For each $h \in \{1, \dots, H\}$, the agent observes the current state $\mathbf{s}^h \in \mathcal{S}$, selects an action $\mathbf{a}^h \sim \pi^{(h)}(\cdot | \mathbf{s}^h) \in \mathcal{A}$, receives a reward $r_h = \mathcal{R}(\mathbf{s}^h, \mathbf{a}^h) \in [0, 1]$, and transitions to the next state $\mathbf{s}^{h+1} \sim \mathcal{T}(\cdot | \mathbf{s}^h, \mathbf{a}^h) \in \mathcal{S}$. We let $\tau_k = (\mathbf{s}_{(k)}^1, \mathbf{a}_{(k)}^1, r_{(k)}^1, \dots, \mathbf{s}_{(k)}^H, \mathbf{a}_{(k)}^H, r_{(k)}^H, \mathbf{s}_{(k)}^{H+1})$ be the random variable that represents the trajectory experienced by the agent in the episode k . $\mathbf{H}_k = (\tau_1, \dots, \tau_{k-1}) \in \mathcal{H}_k$ is the random variable representing the entire history of the agent's interaction with the environment at the start of the episode k . The regret of the RL algorithm over K episodes is: $r_g(K, \pi^{(1)}, \dots, \pi^{(K)}, \mathcal{E}^*) = \bar{V}_{\mathcal{E}^*, 1}^* - \bar{V}_{\mathcal{E}^*, 1}^{\pi^{(k)}}$, with the value function $\bar{V}_{\mathcal{E}, h}^{\pi^{(k)}} = \mathbb{E}_{a \sim \pi^{(k)}(\cdot | \mathbf{s})} [Q_{\mathcal{E}, h}^{\pi^{(k)}}(\mathbf{s}, a)]$. The associated action-value function is $Q_{\mathcal{E}, h}^{\pi^{(k)}}(\mathbf{s}, a) = \mathbb{E} \left[\sum_{h'}^H \mathcal{R}(\mathbf{s}^{h'}, \mathbf{a}^{h'}) \mid \mathbf{s}^h = \mathbf{s}, \mathbf{a}^h = \mathbf{a} \right]$, where the expectation integrates over randomness in the action selections

and transition dynamics. The Bayesian regret is defined as the expected value of the sum of the episodic regrets: $\bar{r}_g(K, \pi^{(1)}, \dots, \pi^{(K)}, \mathcal{E}^*) = \mathbb{E} \left[\sum_{k=1}^K \left[\bar{V}_{\mathcal{E}^*, 1}^* - \bar{V}_{\mathcal{E}^*, 1}^{\pi^{(k)}} \right] \right]$. Note that the regret is a random variable due to the uncertainty in \mathcal{E} . However, we emphasize again here that using a classical MFRL algorithm is not generalizable to multiple wireless environments, which necessitates us revisiting the SC design approach using causal machine learning (ML) tools as explained in later sections. To incorporate causality, the IL framework can be reformulated using an MBRL approach, as proposed in [29]. Compared to MFRL, MBRL is generally more efficient, making it particularly well-suited for SC. This is because MBRL employs a learned model of the environment $p(\widehat{\mathbf{s}}^{t+1} \mid \widehat{\mathbf{s}}^t, \mathbf{a}^t, \boldsymbol{\theta})$ (part of the imitator's job here), aided by a history of received information, to plan and update its policy. However, similar to other supervised learning methods, MBRL faces the challenge of generalization. The data used for training may not match the one encountered during testing, and even small inaccuracies in the dynamics model or changes in the control policy can lead to exploring new parts of the state space. Thus, we need new models that have strong generalization capabilities in MBRL. To achieve the SC goal of improved resource efficiency, we propose to bridge this gap in the literature for IL by developing generalizable (across several wireless environments) solutions that are causality aware.

The mismatch in knowledge between expert agent (\mathcal{E}) and imitator agent ($\widehat{\mathcal{E}}$) that gets learned using MBRL can lead to incorrect inferences due to auto-suggestive delusions [30], i.e., false beliefs generated within one's mind. To address this issue, we present a novel approach by applying a new IL principle, in which we treat actions as causal interventions, as described in [22]. In supervised learning, one can condition or intervene on data using factual and counterfactual error signals, respectively, which helps resolve auto-suggestive delusions. To account for imperfections in the imitator's model, we employ confounding variables in causality. We begin by using SCMs to exploit sparsity in state transitions, which helps generalize to unseen parts of the state-action space. We then extend the IL framework to include some variables $\boldsymbol{\theta} \in \Theta$ observed by the expert but not the imitator. This setup, introduced in [21] for RL, allows modeling of the unknown part of the environment that the imitator learns from its previous state-action sequences. The method in [21] assumes that the imitator has perfect knowledge of \mathbf{s}^t , which differs from our SC system. Also, their AI approach relies on conventional data-driven techniques, such as VAE (without any structure extraction), that may not be sufficient to achieve generalizable wireless systems. Given the MBRL model, in the next step, we discuss how the expert agent computes its expert policy using the state and confounding variables to which it has access. Additionally, we will

explore how a causally invariant semantic language is defined.

2) *Expert policy and causally invariant semantic language*: The source user does not know the confounding variables because it is a computationally constrained node. In this case, the source user communicates the state information to the edge server. Along with confounding variables, the edge server computes an optimal policy, called the *expert policy*, π_E . π_E is computed to maximize the average reward over time, $\pi_E^* = \arg \max_{\pi_E} \mathbb{E}[\mathcal{R}(s)]$, where $\mathcal{R}(s) = \mathbb{E}_{a \sim \pi_E(a^t | s^t, \theta)} \mathcal{R}(s, a)$. π_E is communicated to the source user via a highly reliable backhaul connection. Using π_E , the source user samples transitioned state information. Now, the source user must extract the causality between the state transitions which can be computed as the distribution (see [31] for more details on causal representation as posterior factorization):

$$\mathcal{P}_1 : \text{Compute } p(s^{t+1} | s^t, a^t, \theta) = \prod_i \underbrace{p(s_i^{t+1} | Pa(s_i^t))}_{\text{Causal graph structure}}$$

subject to $a^t = \pi_E^*(a^t | s^t, \theta)$. (1)

The factorization of the transition probability in (1) results from causal discovery. This concept of identifying causal features is closely linked to sparsity, which is a form of *inductive bias* that shapes the behavior of the learning agent. An important aspect of this approach is the continuous improvement of the learning policy, which helps us achieve both causal invariance and sparsity [32]. With this in mind, we can ask whether we can leverage the sparsity of transition dynamics to develop a model that generalizes better to unseen parts of the state-action space. Moreover, using the associational, interventional, and counterfactual levels of reasoning, the source user can generalize to distinct data distributions. Training AI models to generalize over the wireless environment, rather than specializing in fitting specific conditions (e.g., a single cell), provides the necessary robustness to deploy a single AI model for a given task across the entire network. This improves scalability and reduces the complexity (in terms of training + distinct AI models) of the AI functionalities and operations. Additionally, we establish the semantic representation that emerges from the proposed semantic language, which is defined below.

Definition 2. A *semantic language* $\mathcal{L} = (s^t, z^t)$, is a mapping from the observed states s^t , to their corresponding semantic representation z^t , based on the identified SCEs and the causal graph. This mapping is described as the encoder probability distribution $p(z^t | s^t, s^{t+1})$.

To achieve the goal of reducing transmission in the SC system, it is optimal to compute the encoder distribution by determining the NN parameters Ψ that result in the maximum received semantic information at the imitator node, as follows:

$$\begin{aligned} \mathcal{P}_2 : \arg \max_{p_{\Psi}(z^t | (s^t, s^{t+1}))} & \mathbb{I}(s^t; y^t) \\ \text{subject to } \mathbb{I}(s^t; z^t) & \leq \mathbb{I}_b, \end{aligned} \quad (2)$$

where \mathbb{I}_b represents the bandwidth limitations of the wireless channel. The information measures above correspond to the novel semantic information measures detailed in Section III (see equation (7) and (9)). The abstracted state information z^t depends on two factors: 1) the IIT, which helps identify SCEs and rigorously formulate semantic information, and 2) semantic awareness, which is determined by user-defined key performance indicators (KPIs) such as QoE. We next define the QoE metric below.

3) *Semantic effectiveness as QoE metric*: On the imitator's side, the semantic decoder module should be designed to maximize semantic effectiveness defined based on our work [33]. We thus introduce the metric C_t , which captures the causal impact of the expert agent's message (via the imitator's actions) as observed through a channel with a response characterized using $p(y^t | z^t)$ (this distribution could capture the fading and interference in the wireless environment). In other words, C_t measures the semantic effectiveness (inversely proportional to C_t) of the transmitted message to the end-user. We define $C_t(s^t, \hat{s}^t)$ in (3), where $\text{KLD}(p \parallel q)$ represents the KL divergence between p and q . $\pi_{\eta}(a^t | s^t, \hat{\theta}^t)$ is the sub-optimal policy computed by the imitator using the limited information it received over the air. $p(\hat{s}^t | y^t)$ denotes the semantic decoder distribution. The decoder component can thus be designed using the game-theoretic framework we developed in [33]. The second term in the quantity C_t is an average measure across the channel realization. Since the semantic effectiveness metric in (3) captures the channel effect on the transmitted semantics, the learned policy (for the imitator) parameters is robust to the channel errors (in terms of the best the imitator can do). The semantic effectiveness measure is the distance in the semantic space (distance between policy distributions) and not a conventional metric like Euclidean distance. However, computing the semantic effectiveness at the imitator requires it to know π_E . This necessitates communicating the imitator policy to the edge server via the backhaul. Subsequently, the edge server assesses C_t and transmits it to the imitator, which then adapts its policy based on the quality of semantic effectiveness. The communication link between the imitator and the edge server is assumed to be highly reliable compared to the wireless connection to the source user.

Having defined semantic effectiveness, we look at how to define the reward signal mentioned earlier. The agent may encounter extremely sparse or entirely absent extrinsic rewards in numerous real-world situations [34]. In these cases, the intrinsic reward signal based on semantic effectiveness can be used to facilitate collaborative exploration and communication skill learning (specifically, semantic encoder and decoder) between the expert and imitator agents. This skill becomes helpful in subsequent communication

$$C_t(\mathbf{s}^t, \hat{\mathbf{s}}^t) = \text{KLD} \left(\pi_E(\mathbf{a}^t | \mathbf{s}^t, \boldsymbol{\theta}) \parallel \sum_{\mathbf{y}^t, \mathbf{z}^t} p(\mathbf{z}^t | \mathbf{s}^t) p(\mathbf{y}^t | \mathbf{z}^t) p(\hat{\mathbf{s}}^t | \mathbf{y}^t) p(\hat{\boldsymbol{\theta}}^t | \mathbf{H}_t) \pi_\eta(\mathbf{a}^t | \hat{\mathbf{s}}^t, \hat{\boldsymbol{\theta}}^t) \right). \quad (3)$$

phases. Motivated by these considerations, we can define $r_t = r_t^i + r_t^e$, where r_t^e is the extrinsic reward which is mostly (if not always) zero and the intrinsic reward, $r_t^i = \frac{1}{1+C_t(\mathbf{s}^t, \hat{\mathbf{s}}^t)}$. By incorporating semantic effectiveness when computing π_E , the prediction of state transitions and causal structure is improved. This leads to a better design of the semantic encoder on the expert side, which adjusts its transmission strategy to maximize the semantic information extracted by the receiver. Having outlined the different IL components for CSC, we next develop a novel semantic information measure that helps us identify the SCEs and that is critical in formulating various objective functions ($\mathcal{P}_1, \mathcal{P}_2$ in (1) and (2)) elaborated in Section IV.

III. CHARACTERIZING SEMANTIC CONCEPTS VIA INTEGRATED INFORMATION THEORY

An SCM with respect to the formulations in IIT can be described as a stochastic system $\mathcal{U} = \{U_1, \dots, U_n\}$ of n interacting units (atomic) with state space $\Omega_{\mathcal{U}} = \prod_i \Omega_{U_i}$, where $\Omega_{U_i} \subset \mathcal{R}$, and the current state is $\mathbf{u} \in \Omega_{\mathcal{U}}$. Each network state \mathbf{s}_i^t in (1) that is observed at DT can be composed of multiple such atomic units, whose state space we denote as $\Omega_{s_i} \subset \Omega_{\mathcal{U}}$. Our system is updated in discrete steps such that the state space $\Omega_{\mathcal{U}}$ is finite, and the individual random variables $U_i \in \mathcal{U}$ are conditionally independent given the preceding state of \mathcal{U} : $p(\mathbf{u}^{t+1} | \mathbf{u}^t) = \prod_i p(u_i^{t+1} | u^t)$. The IIT-based concept of semantic information draws inspiration from the theory of consciousness measures proposed by [35]. This concept is constructed on the fundamental principles of intrinsic information, information integration, and exclusion, which we elaborate on next.

A. Intrinsic Information for State Abstraction

In IIT, *intrinsic information* refers to the inherent cause-and-effect structure within a system that produces the particular set of observed states and transitions. Put simply, it is the semantic information built into our DT-based system that leads to the shift from one state (\mathbf{s}^t) to the next (\mathbf{s}^{t+1}). Moreover, in IIT, information present in \mathbf{s}^t is considered to be causal only if it has selective causes and selective effects within the system. This selective nature of information distinguishes it from mere correlation or statistical association between variables (which is the extrinsic notion of information used in Shannon's theory of communication), which may not be causal in nature. We now proceed to analytically define intrinsic information. The amount of information that the current state \mathbf{s}_i^t (which represents any subset of causal variables part of the SCM) specifies about the past, i.e., its cause information \mathbb{I}_c , is measured as the distance between the cause repertoire

$p(\mathbf{s}_i^{t-1} | \mathbf{s}_i^t)$ and the unconstrained past repertoire $p(\mathbf{s}_i^{t-1})$, and is defined as follows:

$$\mathbb{I}_c(\mathbf{s}_i^{t-1} | \mathbf{s}_i^t) = \mathbb{D}(p(\mathbf{s}_i^{t-1} | \mathbf{s}_i^t) || p(\mathbf{s}_i^{t-1})). \quad (4)$$

Just like cause information \mathbb{I}_c , the effect information \mathbb{I}_e of \mathbf{s}_i^t is quantified as the distance between the effect repertoire of \mathbf{s}_i^t and the unconstrained future repertoire $p(\mathbf{s}_i^{t+1})$ and written as

$$\mathbb{I}_e(\mathbf{s}_i^{t+1} | \mathbf{s}_i^t) = \mathbb{D}(p(\mathbf{s}_i^{t+1} | \mathbf{s}_i^t) || p(\mathbf{s}_i^{t+1})). \quad (5)$$

Next, we define an appropriate measure for \mathbb{D} . The Kullback-Leibler divergence (KLD) is a useful metric from classical information theory, but it is not a true metric (it is not symmetric) and it is unbounded. Additionally, KLD only measures how “sharp” a distribution is compared to another, without taking into account whether some states of the system are closer than others (in the sense of the Euclidean distance metric). Compared to the literature in IIT [23], a more appropriate measure that aligns better with the IIT notion of information as “differences that make a difference” is the *earth mover's distance (EMD)* or *Wasserstein distance*. This distance is formulated in the context of optimal transport problems as follows:

Definition 3. For any two probability measures μ_s, μ_t , the optimal cost of transporting from μ_s to μ_t can be formulated as the *Wasserstein distance* $W_p^p(\mu_s, \mu_t) = \min_{\gamma \in \mathcal{P}_{\Omega_x \times \Omega_y}} \int \|x - y\|^p \gamma(x, y) dx dy$, where Ω_x and Ω_y are the domain of μ_s and μ_t , respectively. \mathcal{P} is the set of joint probability distributions.

$W_p^p(\mu_s, \mu_t)$ is symmetric and bounded. Finally, having calculated \mathbb{I}_c and \mathbb{I}_e , the total amount of cause-effect information \mathbb{I}_{ce} specified by \mathbf{s}_i^t over the purview defined as $\{\mathbf{s}_i^{t+1}, \mathbf{s}_i^{t-1}\}$, is:

$$\mathbb{I}_{ce}(\mathbf{s}_i^{t+1}, \mathbf{s}_i^{t-1} | \mathbf{s}_i^t) = \min(\mathbb{I}_c, \mathbb{I}_e). \quad (6)$$

From the intrinsic information perspective, each causal mechanism (includes the causes and effects of any state) in the system acts as an information bottleneck. This means that its cause information only exists for the system to the extent that it also specifies effect information, and vice versa.

Lemma 1. We define the intrinsic information learned by the expert and imitator as \mathbb{I}_{ce}^E and \mathbb{I}_{ce}^n , respectively. In the presence of confounding variables $\boldsymbol{\theta}$, the imitator learns less intrinsic information, $\mathbb{I}_{ce}^n < \mathbb{I}_{ce}^E$, compared to the expert. Under the assumption of perfect \mathbf{s}^t being extracted by the imitator, the difference $|\mathbb{I}_{ce}^n - \mathbb{I}_{ce}^E|$ can be shown to be strictly smaller than ϵ when $D_{TV}(p(\boldsymbol{\theta} | \mathbf{H}_t), \delta_{\boldsymbol{\theta}}(\boldsymbol{\theta}^0)) < \epsilon$, where D_{TV} is the total variation distance (see Appendix A) and $\delta_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ is the Dirac-delta distribution, with $\boldsymbol{\theta}^0$ the true value.

Proof: See Appendix B. ■

Lemma 1 emphasizes the significance of precise modeling of the confounding variables (known at the DT side) at the imitator to extract the relevant semantics. Moreover, we introduce the following concept, which enables the DT-enabled transmitter to generate a semantic representation that captures only the relevant and distinctive causal states with nonzero intrinsic information.

Definition 4. We define the *causal-invariant state abstraction* as the D -dimensional embedding of the states, $\phi_i : s \rightarrow \mathcal{R}^D$. ϕ_i is causal-invariant if $\forall s_c, s_e, s_1, s_2 \in \mathcal{S}$, $a \in \mathcal{A}$, $\phi_i(s_1) = \phi_i(s_2)$ if and only if the cause and effect set are same for both s_1, s_2 while satisfying $P(s_1 | s_c, a) = P(s_2 | s_c, a)$, $P(s_e | s_1, a) = P(s_e | s_2, a)$.

From Definition 4 we can conclude the following: if $\phi_i(s_1) = \phi_i(s_2)$, then the intrinsic information conveyed by s_1 and s_2 is the same. In other words, from an SC perspective, state abstraction enables a unique representation for states that share comparable cause-and-effect repertoires. This leads to a reduction in the amount of semantic information conveyed compared to systems that do not employ state abstraction. The concept of state abstraction presented here differs significantly from the traditional ML definitions such as those in [32] and [36]. Next, using the intrinsic information concept, we analyze the information content provided by a state s_i^t composed of several SCEs $U_j \in \mathcal{U}$.

B. Information Integration (via Compositionality, for Identifying Semantic Content Elements)

At the level of an individual causal mechanism (state + its cause and effect mechanisms), the integration postulate [35] states that only mechanisms that specify integrated information can contribute to consciousness. Inspired by this, *integrated information* for an SCM is information that is generated by the whole mechanism beyond the information generated by its parts, meaning that the mechanism is irreducible with respect to information. Similar to cause-effect information, integrated information (denoted by \mathbb{I}_ϕ) is calculated as the Wasserstein distance between two probability distributions.

To compute the quantity of information that is integrated across the m parts of a system, we can partition the system into m parts such as $M_1^t, M_2^t, \dots, M_m^t$. This partition $p_k \in \mathcal{P}_S$ (the set of all partitions of \mathcal{S}) is defined such that $\cup_i M_i^t = s_i^t$ and $M_i^t \cap M_j^t = \emptyset$. The measure of integrated information with respect to the cause and effect mechanisms can be expressed as (7). The integrated information for partition p_k will be:

$$\mathbb{I}_\phi^{p_k} = \min(\mathbb{I}_{\phi,c}^{p_k}, \mathbb{I}_{\phi,e}^{p_k}). \quad (8)$$

\mathbb{I}_ϕ satisfies $0 \leq \mathbb{I}_\phi \leq \min(\mathbb{I}_c(s_i^{t-1}; s_i^t), \mathbb{I}_e(s_i^{t+1}; s_i^t))$. Thus, the SCM is reducible if at least one partition $p_k \in \mathcal{P}_S$ makes no difference to the cause or effect probability, i.e., $\mathbb{I}_\phi^{p_k} = 0$. If it is zero, then the partition does not contribute

any shared information to the system as a whole. Next, we define the integrated information of an SCM as given by its irreducibility over its minimum partition $p_k \in \mathcal{P}_S$:

$$\begin{aligned} \mathbb{I}_\phi &= \mathbb{I}_\phi^{p_k^*}, \\ \text{s.t. } p_k^* &= \arg \min_{p_k} \frac{\mathbb{I}_\phi^{p_k}}{\max_{p_i \in \mathcal{P}_S} \mathbb{I}_\phi^{p_i}}. \end{aligned} \quad (9)$$

The normalization above is over the maximum possible value that $\mathbb{I}_\phi^{p_k}$ could take for any partition.

Having defined the integrated information measure, we consider the IL setup. The imitator relies on the probability distributions obtained from the received semantic information to measure its \mathbb{I}_ϕ . To extract the maximum information, the decoding must be performed optimally using the “true” conditional distribution, which is also a function of confounding variables.

$$p(s_i^t, \theta | s_i^{t-1}) = p(M_1^t, \dots, M_m^t, \theta | M_1^{t-1}, \dots, M_m^{t-1}). \quad (10)$$

To decode s_i^{t-1} , the imitator uses a “false” conditional distribution, $q(\hat{s}_i^t, \hat{\theta} | s_i^{t-1})$ (false, since θ is unknown at the imitator), and hence it is an instance of “mismatched” decoding. \hat{s}_i^t is the received semantic information. To quantify integrated information, we specifically consider the mismatched decoding that uses the “partitioned” probability distribution $q(\hat{s}_i^t, \hat{\theta} | s_i^{t-1}) = p_\eta(\theta | H_t)p(s_i^t | s_i^{t-1}, \theta)$. Similarly, we can define the probability distribution $q(\hat{s}_i^t, \hat{\theta} | s_i^{t+1})$. The imitator should learn (see Section IV) these cause and effect probabilities to extract the semantic information (not just look to reconstruct the current state s_i^t from \mathbf{y}^t). The corresponding learned integrated information, defined as \mathbb{I}_ϕ^η , is called as the extrinsic information (measured using the imitator’s observations, which are the received signals). Next, we quantitatively analyze the error in \mathbb{I}_ϕ between expert and imitator agents for our DT-based SC system.

Lemma 2. *The error between the true integrated information and the extrinsic information learned by the imitator can be quantified as the bias of the learned estimator of cause and effect entropy at the imitator. Here, \mathbb{H} refers to Shannon’s entropy, which is a measure of uncertainty.*

$$\begin{aligned} \mathbb{I}_{\phi,c} - \mathbb{I}_{\phi,c}^\eta &= \mathbb{E}_q (\mathbb{H}(s_i^{t-1} | s_i^t, \theta) - \mathbb{H}(s_i^{t-1} | s_i^t, \theta)), \\ \mathbb{I}_{\phi,e} - \mathbb{I}_{\phi,e}^\eta &= \mathbb{E}_q (\mathbb{H}(s_i^{t+1} | s_i^t, \theta) - \mathbb{H}(s_i^{t+1} | s_i^t, \theta)). \end{aligned} \quad (11)$$

Proof: See Appendix C. ■

Lemma 2 means that the error in integrated information is the minimum of the estimator (of entropy, with respect to the mismatched probability distribution q) bias among cause and effect transitions. In other words, the error in semantic information at the imitator is the same as the error in the transition $(\mathcal{T}, \text{part of } \mathcal{E})$ modeling.

C. Semantic Concepts using Exclusion

A maximally irreducible cause-effect repertoire (MICE) is specified by a subset of elements, referred to as a

$$\mathbb{I}_{\phi,c}^{p_k} = \mathbb{I}_c(\mathbf{s}_i^{t-1}; \mathbf{s}_i^t) - \sum_j \mathbb{I}_c(\mathbf{M}_j^{t-1}; \mathbf{M}_j^t), \quad \mathbb{I}_{\phi,e}^{p_k} = \mathbb{I}_e(\mathbf{s}_i^{t+1}; \mathbf{s}_i^t) - \sum_j \mathbb{I}_e(\mathbf{M}_j^{t+1}; \mathbf{M}_j^t). \quad (7)$$

concept. To find the core cause of any state \mathbf{s}_i^t , we compute $\mathbb{I}_c(\mathbf{M}_i^{t-1} | \mathbf{s}_i^t)$ for all possible partitions \mathbf{M}_i , and take the maximum among them. That is, $\mathbb{I}_c^{\max}(\mathbf{s}_i^t) = \max_{\mathbf{M}_i} \mathbb{I}_c(\mathbf{M}_i^{t-1} | \mathbf{s}_i^t)$. Similarly, we can compute the maximum $\mathbb{I}_e^{\max}(\mathbf{s}_i^t)$ in the effect direction. The partitions that represent the maximum of those values are $\mathbf{m}_c^* = \arg \max_{\mathbf{M}_i} \mathbb{I}_c(\mathbf{M}_i^{t-1} | \mathbf{s}_i^t)$ and $\mathbf{m}_e^* = \arg \max_{\mathbf{M}_i} \mathbb{I}_e(\mathbf{M}_i^{t-1} | \mathbf{s}_i^t)$. The core cause and effect of mechanism \mathbf{s}_i^t are \mathbf{m}_c and \mathbf{m}_e , respectively. Together, they specify the “what” of the concept of \mathbf{s}_i^t . A mechanism that specifies MICE constitutes a concept, or more specifically, a core concept. A concept \mathbf{s}_i^t is composed of several atomic units U_i , referred as *SCEs*.

Definition 5. A *semantic content element* can be formally defined as an atomic mechanism, with possible minimum integrated information among all partitions p_i .

1) *Sub-Optimal Strategy for SCE Identification:* To correctly compute (7), it is necessary to know the true transition probabilities that are used in the evaluation of intrinsic information. However, since the true transition dynamics are unknown, we must rely on estimating an empirical distribution from the available data. As a result, the approach outlined here represents a suboptimal scheme for identifying the SCEs. We define the input data as a sequence of entities, each of which is represented by a d -dimensional value, with each entity being $e_i \in \mathcal{U}$. To extract the entities, we can utilize the β -VAE [37] which is a promising approach (the discussion of which is beyond the scope here) to extract statistically independent entities present in the data. We store the entities for further processing. Further, we define a function f_ϕ to compute the integrated information for a given subset of entities, $\mathbf{u} \subseteq \mathcal{U}$. The function should take the subset of entities \mathbf{u} as input and return the integrated information, $\mathbb{I}_{\phi,\mathbf{u}}$ as in (8). To calculate the integrated information, the procedure utilizes the analytical approach outlined in Section III-B. First, we define a function to generate all possible partitions of a given subset of entities. The function should take the subset of entities as input and return a list of all possible partitions. Further, we iterate over all partitions and compute the integrated information for each partition as in (7) by calling the function f_ϕ . This returns the identified SCEs, which consist of subsets of entities with non-zero integrated information and their corresponding partitions. Among the partitions with nonzero integrated information $\mathbb{I}_{\phi,\mathbf{u}}^{p_k}$, we compute the maximally irreducible partition \mathbf{s}_i^t . The concept as captured by \mathbf{s}_i^t , and its constituent atomic units, U_i , form the SCEs. The computation of the integrated information and generation of partitions can be computationally expensive for large datasets, so it may be necessary

to optimize these functions for performance.

The amount of integrated information generated by concept \mathbf{s}^t is the minimum between past and future:

$$\mathbb{I}_\phi^{\max}(\mathbf{s}_i^t) = \min(\mathbb{I}_c^{\max}(\mathbf{m}_c^* | \mathbf{s}_i^t), \mathbb{I}_e^{\max}(\mathbf{m}_e^* | \mathbf{s}_i^t)). \quad (12)$$

The set of all concepts within an SCM constitutes its conceptual structure, which can be represented in concept space. Concept space is a high-dimensional space, with one axis for each possible past and future state of the system. In this space, each concept is represented as a “star,” with its coordinates given by the probability of past and future states in its cause-effect repertoire. The size of the star represents the value of its $\mathbb{I}_\phi^{\max}(\mathbf{s}_i^t)$. If $\mathbb{I}_\phi^{\max}(\mathbf{s}_i^t)$ is zero, the concept does not exist and if it’s small, the concept exist to a minimal extent. The observed state \mathbf{s}^t at the DT-based expert agent comprises various semantic concepts, each designated as \mathbf{s}_i^t . Even though previous literature, e.g., [38], has briefly mentioned the data observed at the transmitter as composed of multiple concepts, purely as a probabilistic transformation, the proposed rigorous formulation based on theory of consciousness described in this section is the first of its kind. Up to this point, we have identified the SCEs present in the data and analyzed the information they contain. Our next step is to comprehend the associations between the different SCEs and how they collectively form a semantic concept, ultimately creating a causal graph that gives structure to the data.

D. Defining Causal Relationships via Topological Perspective

For simplicity, we remove the dependence on t in this subsection. Any concept $\mathbf{m} \subseteq \mathbf{s}$ with $\mathbb{I}_\phi(\mathbf{m}) > 0$ specifies a candidate distinction $d(\mathbf{m}) = (\mathbf{m}, \mathbf{z}^*, \mathbb{I}_\phi(\mathbf{m}))$ within the system \mathcal{S} in state \mathbf{s} . $\mathbf{z}^* = \{\mathbf{z}_c^*, \mathbf{z}_e^*\}$ represents the maximal cause-effect pair for \mathbf{m} . For any state \mathbf{s} , the causal distinctions that represent its subsets with nonzero integrated information are defined as (13), where $d(\mathbf{m})$ is the set of nonzero integrated information for \mathbf{m} and $\mathbf{s}' = \{\mathbf{s}_c', \mathbf{s}_e'\}$ is the maximal cause-effect pair for the system as a whole. *Causal relationships* refer to the way in which the causes and/or effects of a set of distinctions $\mathcal{D}(\mathcal{T}_S, \mathbf{s})$ in a complex system overlap. In the same way that a distinction \mathbf{m} identifies which units/states constitute a cause and the resulting effect, a relationship identifies which units/states correspond to which units/states among the purviews of a set of distinctions. These relationships reflect how the cause-and-effect power of its concepts is interconnected within the complex system. In short, they help to attribute *meaning (semantics)* to the state \mathbf{s} . Understanding causal relationships is crucial, as these models enable DT-based wireless systems to distill knowledge and experiences,

$$\mathcal{D}(\mathcal{T}_S, \mathbf{s}) = \{d(\mathbf{m}) : \mathbf{m} \subseteq \mathbf{s}, \mathbb{I}_\phi(\mathbf{m}) > 0, \mathbf{z}_c(\mathbf{m}) \subseteq \mathbf{s}'_c, \mathbf{z}_e(\mathbf{m}) \subseteq \mathbf{s}'_e\}, \quad (13)$$

similar to the human mind, to make reliable predictions, generalizations, and analogies. These qualities are essential for enabling compositional and counterfactual reasoning, active intervention in the world to test hypotheses, and the ability to articulate one's understanding to others. We now look at how to quantify the causal relationships for our system. The degree of irreducibility resulting from this binding of cause-and-effect power is measured by the irreducibility of the relationships, denoted by $(\mathbb{I}_{\phi_r} > 0)$. In the presence of confounding variables, the learned causal distinctions by the imitator follows the relation, $\mathcal{D}(\mathcal{T}_S, \mathbf{s} | \hat{\theta}) \subseteq \mathcal{D}(\mathcal{T}_S, \mathbf{s} | \theta)$. As in [24], the relations between any two distinctions can be defined as follows. For any state \mathbf{s} , consider two causal distinctions $d(\mathbf{m}_1)$ and $d(\mathbf{m}_2)$, that have a possibly non-empty overlapping atomic units as their constituents. The relation between $d(\mathbf{m}_1)$ and $d(\mathbf{m}_2)$ can be defined as the maximally irreducible subset that is common to both of them. For each candidate overlap, the intrinsic difference is assessed at the maximum information partition (MIP) for each partition, and summed together, which represents the irreducibility measure. If we consider a specific set of distinctions $\mathbf{d} \subseteq \mathcal{D}(\mathcal{T}_S, \mathbf{s})$, it is possible that there are multiple sets of causes and/or effects, denoted as \mathbf{z} , such that

$$z : z \cap \{z_c^*(d), z_e^*(d)\} \neq \emptyset \forall d \in \mathbf{d}, \cap_{z \in \mathbf{z}} z \neq \emptyset, |\mathbf{z}| > 1, \quad (14)$$

with maximal overlap (called the “faces” of the relation), $o^*(z) = \cap_{z \in \mathbf{z}} z \neq \emptyset$. A relation $r(\mathbf{d})$ thus consists of a set of distinctions $\mathbf{d} \subseteq \mathcal{D}(\mathcal{T}_S, \mathbf{s})$, with an associated set of faces $f(\mathbf{d}) = \{f(z)\}_{\mathbf{d}}$ and irreducibility measure $\phi_r > 0$ and can be written as $w(\mathbf{d}) = (\mathbf{d}, f(\mathbf{d}), \phi_r)$. $w(\mathbf{d})$ can be topologically represented as shown below.

Lemma 3. *The structure of causes and effects in an SCM can be represented using an abstract simplicial complex [39], in which causes and effects are the vertices and the relations are simplices; relations between pairs of causes and effects are represented as 2-simplices (edges); relations between trios are represented as 3-simplices (faces) and so on.*

The Lemma 3, whose proof follows directly from its definition states that a relation $r(\mathbf{d})$ can be topologically characterized as a $|\mathbf{d}|$ -simplex, where $|\mathbf{d}|$ is the order of a relation. Relations between two purviews are known as 2-relations, relations involving three purviews are referred to as 3-relations and so on. More details on simplex is provided in Appendix A. The degree of a relation is defined as the number of causal distinctions present in \mathbf{d} . The “faces” of the relation is same as the concept of faces in simplicial homology. Further, we look at defining the irreducibility measure ϕ_r . For any relation $w \in \mathcal{W}$, the irreducibility can be defined as

$$\phi(\mathcal{R} | \mathcal{O}) = \sum_i d_{\mathcal{O}}(\mathcal{S}_i), \quad (15)$$

where \mathcal{O} represents a candidate overlap between two causal distinctions, \mathcal{S}_i represents a particular partition of \mathcal{O} and $d_{\mathcal{O}}(\mathcal{S}_i) = D(p(\mathcal{O}) \parallel p(\mathcal{S}_i))$. Hence, $d_{\mathcal{O}}(\mathcal{S}_i)$ represents the information the partition \mathcal{S}_i provides within a specific overlap \mathcal{O} . Finally, the maximally irreducible overlap is the \mathcal{O} that maximizes ϕ :

$$\begin{aligned} \phi_{\max} &= \arg \max_{\mathcal{O}} \phi(\mathcal{S} | \mathcal{O}). \\ \phi_R &= \phi(\mathcal{S} | \mathcal{O}_{\max}). \end{aligned} \quad (16)$$

We now formally define semantics using the abstract simplicial complex-based definition of causal relations.

1) *Topological characterization of semantics:* Semantics are represented using the relations defined in the previous subsection, as follows:

Theorem 1. *An abstract cell complex $\mathcal{C} = (\mathcal{D}, \prec_w, \dim)$ is a set \mathcal{D} of abstract SCEs, which are the cells here, equipped with a bounding relation \prec_w and a dimension function assigning to each $\mathbf{d} \in \mathcal{D}$ a non-negative integer (\dim), satisfying the following properties*

- if $\mathbf{d}_1 \prec_w \mathbf{d}_2$ and $\mathbf{d}_2 \prec_w \mathbf{d}_3$, then $\mathbf{d}_1 \prec_w \mathbf{d}_3$ (called as the transitivity property)
- if $\mathbf{d}_1 \prec_w \mathbf{d}_2$, then $\dim(\mathbf{d}_1) < \dim(\mathbf{d}_2)$ (called as the monotonicity property).

Proof: The proof is given in Appendix D. ■

Further, we look at what Theorem 1 entails for the DT-based SC system. $\mathbf{d} \in \mathcal{D}$ represents a causal mechanism (concept) with non-zero integrated information. The dimension of \mathbf{d} , denoted $\dim(\mathbf{d})$, represents the number of distinct SCEs part of \mathbf{d} . The relation between two causal mechanisms \mathbf{d}_1 and \mathbf{d}_2 can be defined as the causal relationship (the overlap in the causal distinctions) between them. The usefulness of the Theorem 1 for our DT-based SC system can be two-fold: Firstly, just like in natural language, where syntax refers to the grammatical forms used to express content and semantics refers to the meaning attributed to those syntactic expressions, in communication between intelligent agents, the encoded representation refers to the syntax part [33]. The transmit encoders, which form a component of the semantic language used to convey meaning, can be connected to the specific problem formulation. The design of the AI architecture employed for this purpose may be constrained by the complexity of the transmitter, such as the maximum number of parameters or layers it can support. Regardless of this dependency on language, if the receiver can extract the semantic content defined using a topological construct such as a cell complex, it should be able to communicate seamlessly with multiple transmitters without requiring any protocol changes. In other words, regardless of the

encoded representations used to convey the states observed at DT, the topological characterization of semantics acts as a “bridge” [40] that unifies different semantic contents and facilitates communication between intelligent agents. The abstract cell complex structures in Theorem 1 can be seen as “universal translators” and bridges across different knowledge representations. Additionally, the topological characterization presented in Theorem 1 allows a rigorous formulation of semantic metrics, such as similarity and reliability, as defined in Section IV-A, compared to the current state of the art. Having defined SCEs present in the observed states, we further move on to discovering the causal structure in the data using novel AI tools.

IV. CAUSAL DISCOVERY AND INFERENCE VIA IMITATION LEARNING AND MODEL BASED REINFORCEMENT LEARNING

In the context of DT-based SC, the causal dynamics model $p(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta})$ cannot be accessed by the imitator due to hidden confounding variables (known only to DT). Instead, the imitator can only estimate the transition distribution based on its interactions with the expert. To enable the imitator to make better control decisions, it must have the ability to develop representations of the world based on past experiences (received semantics from the DT node), which allow for generalization to novel situations. To achieve this, the transmitter design must incorporate causal discovery, state transition probability, and semantic encoding such that the QoE is maximized (i.e., expected imagined rewards $\mathbb{E}(\sum_{\tau=t}^H \gamma^{\tau-t} r_\tau)$), as discussed next.

A. Causal Imitation Learning Problem: Transmitter Design

Herein, at the expert node, our objective is to throw away irrelevant state variables while learning the causal dynamics and hence design an abstract causal state description (represented using $p(\mathbf{z}^t | \mathbf{s}^t, \mathbf{s}^{t+1})$) which gets transmitted. Towards this objectives, we adopt the concept of deep end-to-end causal inference (DECI) from [28]. We specifically selected DECI because other contemporary works treat causal discovery and causal inference as distinct problems. In contrast, DECI is the only approach that simultaneously addresses both problems by conducting causal discovery through the estimation of a sparse directed graph and performing causal inference by estimating specific quantities based on a given set of inputs (interventions) – such as estimating \mathbf{s}^{t+1} in this particular case. Note that the objectives for causal inference in an SC system differ from those in [28]. In contrast to [28], we suggest a constrained optimization approach to achieve a minimum QoE. DECI employs a Bayesian perspective in its causal discovery process, where the causal graph \mathbf{G} is modeled in conjunction with the observations $\mathbf{s}^t, \mathbf{s}^{t+1}$. First, we look at the joint distribution (17), where \mathbf{G} is defined as the adjacency matrix (corresponding to the graph \mathcal{G}), with

entries $G_{j,i} \in \{0, 1\}$ indicating the presence of an edge from $j \rightarrow i$. $\boldsymbol{\Psi}$ is the set of NN parameters. The causal relationships can be modeled using a non-linear additive noise model (ANM) as follows: $\mathbf{s}_i^{t+1} = \mathbf{f}_i(\mathbf{s}_i^t, \mathbf{P}a_i) + \mathbf{n}_i$, where \mathbf{n}_i is an exogenous noise variable that is independent of all other variables. We propose a flexible NN parameterization that satisfies the graph adjacency constraints by setting $\mathbf{f}_i(\mathbf{s}_i^{t+1}) = \zeta_i \left(\sum_j G_{j,i} \kappa_j(\mathbf{s}_j^t) \right)$, where the each entry of ζ_i and κ_j are multilayer perceptrons (MLPs). Our objective is to use observational data (that involves network state \mathbf{s}^t) to fit the parameters $\boldsymbol{\Psi}$ of our non-linear ANM. After fitting the model, the posterior $p_{\boldsymbol{\Psi}}(\mathbf{G} | \mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{a}^t, \boldsymbol{\theta})$ reflects our understanding of the causal structure. The graph prior $p(\mathbf{G})$ should characterize the graph as a *directed acyclic graph (DAG)*. We choose it as $p(\mathbf{G}) \propto \exp(-\lambda_s \|\mathbf{G}\|^2 - \rho h(\mathbf{G})^2 - \alpha h(\mathbf{G}))$, where the DAG penalty $h(\mathbf{G}) = \text{tr}(e^{\mathbf{G} \odot \mathbf{G}}) - D$, which is non-negative and zero only if \mathbf{G} is a DAG. Variables α and ρ represent the weights for the DAG penalty [41], and they are optimized during training as in [28, Appendix B]. The prior knowledge about graph sparseness is modeled by the term $\lambda_s \|\mathbf{G}\|^2$, where λ_s is a fixed known quantity. The model has two difficulties: first, the actual posterior over \mathbf{G} cannot be computed efficiently, and second, the maximum likelihood approach cannot be applied to estimate the model parameters due to the presence of the latent variable \mathbf{G} . To address both challenges simultaneously, we employ variational inference, as suggested in previous literature [42]. Specifically, we introduce a variational distribution $q_{\boldsymbol{\Psi}}(\mathbf{G})$ to approximate the intractable posterior $p_{\boldsymbol{\Psi}}(\mathbf{G} | \mathbf{z}^t, \mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{a}^t, \boldsymbol{\theta})$, and utilize it to construct the evidence lower bound (ELBO) given by (18). Here χ is the NN parameters for the $p_{\chi}(x^{(n)} | \mathbf{G})$. (18) can be rewritten as (19). The causal discovery architecture includes a graph NN (GNN) [43] which accepts as input \mathbf{s}^t and propagates information across a fully connected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$. This graph includes vertices $v_i \in \mathcal{V}$ for each \mathbf{s}_i^t , and each pair of vertices (v_i, v_j) is connected by an edge whose embedding is represented by the MLP $e_{i,j} = \mathbf{f}_{i,j}(\mathbf{s}_i^t, \mathbf{s}_j^t)$, and:

$$p_{\boldsymbol{\Psi}}(G_{j,i} | \mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{a}^t, \boldsymbol{\theta}) = \text{Softmax}(e_{i,j} | T), \quad (20)$$

where T represents the softmax temperature variable [44]. DECI aims to demonstrate that the optimization of the variational objective in equation (18) can recover the ground truth data generation mechanism. To establish this statistical guarantee under the correct specification of DECI, which assumes the existence of χ^0 and \mathbf{G}^0 that correspond to the true data generating process and the absence of unobserved confounding variables, two technical assumptions must be satisfied. Firstly, function f_i must be non-invertible (since a DAG), 3rd-order differentiable, and not constant with respect to any of its inputs. Secondly, the proper noise densities must have bounded likelihood. These assumptions

$$p_{\Psi}(\mathbf{z}^t, \mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{a}^t, \boldsymbol{\theta}, G) = p(G)p(\boldsymbol{\theta})p(\mathbf{s}^t)\pi_E(\mathbf{a}^t \mid \mathbf{s}^t, \boldsymbol{\theta})p(\mathbf{s}^{t+1} \mid \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta})p(\mathbf{z}^t \mid \mathbf{s}^t, \mathbf{s}^{t+1}), \quad (17)$$

$$g_E(\boldsymbol{\chi}, \boldsymbol{\Psi}) = \mathbb{E}_{q_{\Psi}(\mathbf{G})} \left[\log p((\mathbf{G}) \prod_{n: x^{(n)} \in \{\mathbf{z}^t, \mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{a}^t, \boldsymbol{\theta}\}} p_{\boldsymbol{\chi}}(x^{(n)} \mid \mathbf{G}) \right] + \mathbb{H}(q_{\Psi}) \leq \log p_{\boldsymbol{\chi}}(x^{(1)}, \dots, x^{(n)}). \quad (18)$$

$$\begin{aligned} g_E(\boldsymbol{\chi}, \boldsymbol{\Psi}) &= \mathbb{E}_{q_{\Psi}(\mathbf{G})} \left[\prod_{n: x^{(n)} \in \{\mathbf{z}^t, \mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{a}^t, \boldsymbol{\theta}\}} p_{\boldsymbol{\chi}}(x^{(n)} \mid \mathbf{G}) \right] - \text{KLD}(q_{\Psi}(\mathbf{G} \mid \mathbf{x}) \parallel p(\mathbf{G})) \\ &\leq \log p_{\boldsymbol{\chi}}(x^{(1)}, \dots, x^{(n)}). \end{aligned} \quad (19)$$

rule out non-identifiable ANM. Further, the ELBO maximization problem \mathcal{P}_1 in (1) that involves expert policy computation, causal discovery and state transition dynamics can be reformulated as follows.

$$\begin{aligned} \mathcal{P}_1 : \quad & \arg \max_{\pi_E(\mathbf{a}^t \mid \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta}), \mathbf{G}, p(\mathbf{s}^{t+1} \mid \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta})} g_E(\boldsymbol{\chi}, \boldsymbol{\Psi}) \\ \text{subject to } & \mathbb{E}(r_t) \leq R_b \end{aligned} \quad (21)$$

We now explain how to obtain the reward constraint bound, denoted as R_b in this context. To accomplish this, we must first define the concept of “semantic reliability”. The expert agent’s goal is to select an optimal semantic representation (encoder) that accurately represents the causal structure at the transmit side. The imitator, on the other hand, aims to improve semantic effectiveness on its side to achieve a desired level of semantic reliability. *Semantic reliability* is quantified by the expression, $p(E_t(\mathbf{s}^t, \hat{\mathbf{s}}^t) < \delta) \geq 1 - \epsilon$, where $E_t(\mathbf{s}^t, \hat{\mathbf{s}}^t) = \|\mathbf{s}^t - \hat{\mathbf{s}}^t\|^2$ represents semantic distortion and ϵ is arbitrarily small. This metric reflects the imitator’s ability to reliably reconstruct all the causal aspects in the decoded causal structure. Unlike classical reliability measures, in semantics, we can recover the actual meaning of transmitted messages even with a higher bit error rate (BER), as long as the semantic distortion remains within the set limit. This is illustrated by the choice of δ here, which depends on the concept of semantic space. The *semantic space* \mathcal{K} is defined as an N -dimensional topological ball (since \mathbf{s}^t belongs to an abstract cell complex as proved in Theorem 1) centered at the actual state \mathbf{s}^t , where all points inside the ball corresponding to the same semantic information. Formally, we can express this as $E_t(\mathbf{s}^t, \hat{\mathbf{s}}^t) \leq \delta$ such that $\mathbb{I}_{\phi}(\mathbf{s}^t) = \mathbb{I}_{\phi}(\hat{\mathbf{s}}^t)$, where E_t is the topological distance between states \mathbf{s}^t and $\hat{\mathbf{s}}^t$, δ is the radius of the topological ball, and \mathbb{I}_{ϕ} is a mapping function that maps states to their corresponding semantic information. The topological distance here refers to the irreducibility measure quantifying the causal relationships between any two states $\mathbf{s}^t, \hat{\mathbf{s}}^t$. The upper bound R_b is defined by the reward obtained at $(\hat{\mathbf{s}}^t)^*$ which corresponds to $(\hat{\mathbf{s}}^t)^* = \arg \min_{\hat{\mathbf{s}}^t \in \mathcal{K}} C_t(\mathbf{s}^t, \hat{\mathbf{s}}^t)$.

The second challenge pertains to resolving the semantic encoder ($p_{\Psi}(\mathbf{z}^t \mid (\mathbf{s}^t, \mathbf{s}^{t+1}))$) problem using a revised version of the information bottleneck principle [45]. In contrast to [45], we leverage semantic information measures

derived from IIT principles. As a result, our metrics go beyond conventional information theory approaches. Our objective is to determine the distribution $p_{\Psi}(\mathbf{z}^t \mid (\mathbf{s}^t, \mathbf{s}^{t+1}))$, which converts the input signal \mathbf{s} into a representation \mathbf{z} , such that \mathbf{z} discloses as little information as possible about \mathbf{s} , while extracting the maximum amount of semantic information about \mathbf{s} in the output signal \mathbf{y} ($\mathbb{I}_{\phi}(\mathbf{s}; \mathbf{y})$). We reformulate \mathcal{P}_2 in (2) as:

$$\begin{aligned} \mathcal{P}_2 : \quad & \arg \max_{p_{\Psi}(\mathbf{z}^t \mid (\mathbf{s}^t, \mathbf{s}^{t+1}))} \mathbb{I}_{\phi}(\mathbf{s}^t; \mathbf{y}^t) \\ \text{s.t. } & \mathbb{I}_{\phi}(\mathbf{s}^t; \mathbf{z}^t) \leq \mathbb{I}_b \end{aligned} \quad (22)$$

According to Definition 4, causally invariant states are represented using the same semantic representation \mathbf{z}^t , which is ensured by the constraint imposed on $\mathbb{I}_{\phi}(\mathbf{s}^t; \mathbf{z}^t)$ in (22). Next, we discuss how to solve the above two optimization problems \mathcal{P}_1 and \mathcal{P}_2 .

1) *Proposed Solution*: The resulting bi-level optimization involving \mathcal{P}_1 and \mathcal{P}_2 can be alternatively solved as follows. We first write the Lagrangian corresponding to (21) and (22) as follows.

$$\begin{aligned} \mathcal{L}_1 = & \min_{\lambda_1} \max_{\pi_E(\mathbf{a}^t \mid \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta}), \mathbf{G}, p(\mathbf{s}^{t+1} \mid \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta})} g_E + \lambda_1 (R_b - \mathbb{E}(r_t)), \\ \mathcal{L}_2 = & \min_{\lambda_2} \max_{p_{\Psi}(\mathbf{z}^t \mid (\mathbf{s}^t, \mathbf{s}^{t+1}))} \mathbb{I}_{\phi}(\mathbf{s}^t; \mathbf{y}^t) + \lambda_1 (\mathbb{I}_b - \mathbb{I}_{\phi}(\mathbf{s}^t; \mathbf{z}^t)). \end{aligned} \quad (23)$$

To solve the resulting bi-level optimization, the technique is described in Appendix G. The resulting alternating updates for outer optimization (solving \mathcal{L}_1) can be obtained as (where (k) denotes iteration) (24). Our approach involves aggressively improving the model to maximize the ELBO based on the current semantic encoder, followed by implementing a more conservative policy and causal discovery. The detailed algorithmic process is presented in Algorithm 1 and the model architecture in Fig. 3, which concludes the SC system design at the expert agent side. The algorithm convergence follows similar theoretical arguments as in [28, Theorem 1]. The authors therein demonstrate that maximizing the ELBO in (18) can recover both the ground truth data generating process, $p(\mathbf{s}^t, \mathbf{s}^{t+1}, \boldsymbol{\theta}, \mathbf{z}^t, \mathbf{a}^t; \mathbf{G}^0)$, and the true causal graph, \mathbf{G}^0 , in the infinite data limit. Apart from the ANM model, the convergence guarantees are under the assumption that 1) there are no latent confounders (which is true from the perspective of expert agent) and 2) the log-likelihood is regular,

$$\begin{aligned}
(\pi_E(\mathbf{a}^t | \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta}))_{k+1} &= \arg \max_{\pi_E(\mathbf{a}^t | \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta})} \mathcal{L}_1 \left(\pi_E(\mathbf{a}^t | \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta}), \mathbf{G}_{(k)}, p_{(k)}(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta}), \lambda_{1,(k)} \right), \\
\mathbf{G}_{(k+1)} &= \arg \max_{\mathbf{G}} \mathcal{L}_1 \left((\pi_E(\mathbf{a}^t | \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta}))_{(k+1)}, \mathbf{G}, p_{(k)}(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta}), \lambda_{1,(k)} \right), \\
p_{(k+1)}(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta}) &= \arg \max_{\pi_E(\mathbf{a}^t | \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta})} \mathcal{L}_1 \left((\pi_E(\mathbf{a}^t | \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta}))_{k+1}, \mathbf{G}_{(k+1)}, p(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta}), \lambda_{1,(k)} \right), \\
\lambda_{(k+1)} &= \arg \min_{\lambda} \mathcal{L}_1 \left((\pi_E(\mathbf{a}^t | \mathbf{s}^{t+1}, \mathbf{s}^t, \boldsymbol{\theta}))_{k+1}, \mathbf{G}_{(k+1)}, p_{(k+1)}(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta}), \lambda_1 \right).
\end{aligned} \tag{24}$$

Algorithm 1 Proposed Bi-level Optimization for Causal Discovery and Semantic Encoder

Given: $p(\mathbf{G}), p(\mathbf{s}^0)$

Define: $\mathcal{E} = (\mathbf{G}, p(\mathbf{s}^{t+1} | \mathbf{s}^t, \boldsymbol{\theta}, \mathbf{a}^t), \pi_E(\mathbf{a}^t | \mathbf{s}^t, \boldsymbol{\theta}))$

Initialize: Policy π_E , causal graph $\mathbf{G} \sim p(\mathbf{G})$, data buffer $\mathcal{D} = \{\}$.

- 1: **for** $(t = 0, 1, 2, \dots)$ **do**
- 2: Collect data $\mathcal{D}_t = (\mathbf{s}^t, \boldsymbol{\theta})$ by executing π_E in the environment.
- 3: Build local (policy-specific) dynamics model: $[\hat{\mathcal{E}}_{t+1}, \hat{\lambda}_{t+1}] = \arg \min_{\lambda} \max_{\mathcal{E}} g_E + \lambda(R_b - \mathbb{E}(r_t))$.
- 4: Improve policy: $(\pi_E)_{k+1} = (\pi_E)_k + \alpha_k \nabla_{\pi_E} g_E(\pi_\eta, \hat{\mathcal{E}}_{k+1}^\pi)$. with a conservative algorithm like policy search with natural gradient (NPG) [46].
- 5: Compute the encoder distribution, given the causal graph and state transitions as $p(\mathbf{z}^t | \mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{G})$ using VAE.
- 6: Simulate the receiver part using Algorithm 2.
- 7: Evaluate semantic effectiveness metric $C_t(\mathbf{s}^t, \hat{\mathbf{s}}^t)$.

8: **end for**

which means that $\mathbb{E}_p [|\log p_\chi(\mathbf{s}^t, \mathbf{s}^{t+1}, \mathbf{z}^t, \mathbf{a}^t | \mathbf{G})|] < \infty$. We next look at the causal IL problem for the receiver design.

B. Causal Imitation Learning Problem: Receiver Design

The distribution of the trajectory τ of the expert's interactions with the environment can be written as:

$$P_E(\tau | \boldsymbol{\theta}) = p(\mathbf{s}^0 | \boldsymbol{\theta}) \prod_{t=0}^H p(\mathbf{s}^{t+1} | \mathbf{s}^t, \mathbf{a}^t, \boldsymbol{\theta}) \pi_E(\mathbf{a}^t | \mathbf{s}^t, \boldsymbol{\theta}) \tag{25}$$

Since the imitator does not have access to the confounding variables $\boldsymbol{\theta}$, it must learn the distribution over trajectories based on the state transitions history that is available to it. This history is constructed from the imitator's interactions with the environment. The imitator can then use this learned distribution to guide its own actions and improve its performance on the task.

$$\begin{aligned}
P_\eta(\tau | \boldsymbol{\theta}) &= \\
p(\hat{\mathbf{s}}^0 | \boldsymbol{\theta}) \prod_{t=0}^H &\underbrace{p(\hat{\mathbf{s}}^{t+1} | \hat{\mathbf{s}}^t, \mathbf{a}^t, \boldsymbol{\theta})}_{\text{causal dynamics model}} \underbrace{\pi_\eta(\mathbf{a}^t | \hat{\mathbf{s}}_0, \mathbf{a}_0, \dots, \hat{\mathbf{s}}^t)}_{\text{imitator policy}}.
\end{aligned} \tag{26}$$

The goal of IL here is for the imitator to compute the policy such that the average reward (which is essentially QoE here) is maximized. The average is taken across all possible trajectories and the confounding variables, i.e., $\mathbb{E}_{\boldsymbol{\theta} \sim P(\boldsymbol{\theta} | \tau)} \mathbb{E}_{\tau \sim P(\tau | \boldsymbol{\theta})} [\mathcal{R}(\hat{\mathbf{s}}^t, \mathbf{a}^t; \boldsymbol{\theta})]$. The resulting policy learned at the imitator is $\pi_\eta(\mathbf{a}^t | \hat{\mathbf{s}}^t, \hat{\boldsymbol{\theta}})$. However, this approach assumes that the imitator is ideal, i.e., it has infinite memory capacity and perceives the state and actions in the same way as indicated by the expert. Specifically, the optimal policy is often a deterministic function of the environment $\mathcal{E} = (\mathcal{T}, \mathcal{R}, \boldsymbol{\theta})$ such that if the imitator is able to identify the environment, it has all the necessary information to determine the optimal policy. The imitator reflects its initial uncertainty about the environment through a prior distribution $P(\mathcal{E} \in \cdot)$. As the history unfolds, the imitator's current knowledge of the environment can be represented by the posterior probabilities $P(\mathcal{E} \in \cdot | \mathbf{H}_t)$. \mathbf{H}_t represents the history of state-action pairs available at the imitator node until time t . The total number of bits needed to identify the environment is $\mathbb{H}(\mathcal{E})$. Assuming the reward function is also unknown to the imitator (note that the expert agent evaluates the semantic effectiveness, which is the intrinsic reward), the total number bits required to learn the environment is

$$\mathbb{H}(\mathcal{E}) = \mathbb{H}(\mathcal{T}) + \mathbb{H}(\mathcal{R}). \tag{27}$$

In other words, the DT-based SC system may need to invest $\mathbb{H}(\mathcal{E})$ bits for the initial training. However, in practice, the imitator can only acquire a learning target $\hat{\mathcal{E}}$, which is a stochastic function of the actual environment. This means that due to communication and computational limitations, the amount of information acquired by the agent about \mathcal{E} is only $\mathbb{I}(\mathcal{E}; \hat{\mathcal{E}})$, which is derived from the semantic information received (states here). This can be quantified as only a subset of history being stored at the agent due to memory limitations, i.e., $P(\mathcal{E} \in \cdot | \mathbf{H}_t^{t-K})$. Here, \mathbf{H}_t^{t-K} represents the history of state transitions corresponding to the most recent K time instants. The states perceived by the imitator can be erroneous due to the presence of a wireless channel between the expert and the imitator. Thus, we can consider that the imitator only knows an estimate $\hat{\mathbf{H}}_t$ and, hence, the environment $P(\hat{\mathcal{E}} \in \cdot | \hat{\mathbf{H}}_t)$. This limits the space over which the learned target environment lies. As a result, aiming for $\hat{\mathcal{E}}$ incurs a bounded degree of performance, $\mathbb{E}[\bar{V}^* - \bar{V}^{\pi_{\hat{\mathcal{E}}}}]$. A natural measure of distortion here is the expected squared regret between the optimal and target policies:

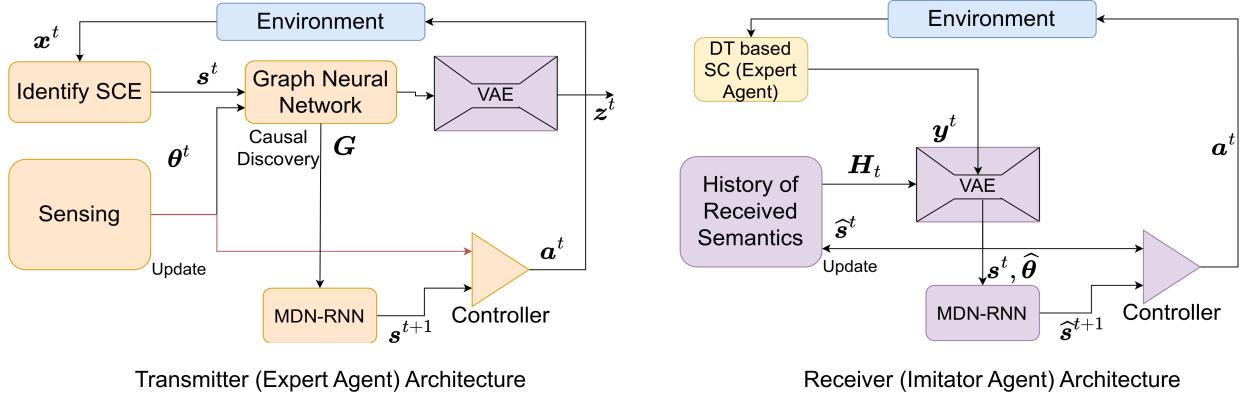


Fig. 3: Generative AI components of the proposed network state model.

$$d(\mathcal{E}, \hat{\mathcal{E}} \mid \mathbf{H}_t) = \mathbb{E} \left[\mathbb{E}_{\pi_{\hat{\mathcal{E}}}} \left(\sum_{t'=t}^H \bar{V}^*(\mathbf{H}_{t'}) - \bar{V}^{\pi_{\hat{\mathcal{E}}}(\mathbf{H}_{t'})} \right)^2 \mid \mathcal{E}, \mathbf{H}_t \right]. \quad (28)$$

We now look at how to estimate the different components of $\hat{\mathcal{E}}$. The receiver begins by extracting the semantics s^t from the received signal y^t , which are modeled using the channel probability distribution $p(y^t \mid z^t)$. The resulting \hat{s}^t is then stored in the observation history. Using this history, the receiver obtains an estimate of the confounding variables $\hat{\theta}^t$ that helped the expert agent to analyze the optimal policy. The imitator then uses both the confounding variable information and the extracted state to derive the policy $\pi_\eta(a^t \mid \hat{s}^t, \hat{\theta}^t)$. We now look at the details of the proposed semantic decoder and imitator policy design at the receiver. First, we define the loss function that captures the difference in learned state transition dynamics compared to that at the expert node.

Definition 6. Given the inferred model at the imitator $\hat{\mathcal{E}}$ and the original sampling distribution $\mu(s, a)$, the *model approximation loss* is defined as follows.

$$l(\hat{\mathcal{E}}, \mu) = \mathbb{E}_{(s, a) \sim \mu} \left[\text{KLD} \left((P(\cdot \mid s, a) \parallel \hat{P}(\cdot \mid s, a)) \right) \right]. \quad (29)$$

(29) is used to quantify the semantic reconstruction loss on the imitator's end. As explained next, we aim to address the issue of determining the state transition probability and imitator policy.

1) *Variational Inference Framework for Receiver Design:* For simplicity, we reuse the same notation Ψ from Section IV-A for the NN parameters here. Our proposed method here learns confounding variables by maximizing an information theoretic objective defined as follows

$$\begin{aligned} F(\Psi) &= \mathbb{H}(s^t; \theta \mid \mathbf{H}_t) + \mathbb{H}[\mathbf{A} \mid s^t; \mathbf{H}_t] - \mathbb{H}(\mathbf{A}; \theta \mid s^t; \mathbf{H}_t) \\ &= (\mathbb{H}[\theta \mid \mathbf{H}_t] - \mathbb{H}[\theta \mid s^t; \mathbf{H}_t]) + \mathbb{H}[\mathbf{A} \mid s^t; \mathbf{H}_t] \\ &\quad - (\mathbb{H}[\mathbf{A} \mid s^t; \mathbf{H}_t] - \mathbb{H}[\mathbf{A} \mid s^t, \theta; \mathbf{H}_t]) \\ &= \mathbb{H}[\theta \mid \mathbf{H}_t] - \mathbb{H}[\theta \mid s^t; \mathbf{H}_t] + \mathbb{H}[\mathbf{A} \mid s^t, \theta; \mathbf{H}_t]. \end{aligned} \quad (30)$$

We particularly derived $F(\Psi)$ (where Ψ captures the NN parameters) for the following reasons. The first term encourages our prior distribution over $p(\theta \mid \mathbf{H}_t)$ to have high entropy, which implies that the imitator ideally prefers to obtain maximum information about θ . The second term suggests that it should be easy to infer the confounding variable θ from the current state and history. The third term suggests that each learned target should act as randomly as possible (random actions), which we achieve by using a maximum entropy policy to represent each confounder. As we cannot integrate over all states and confounders to compute $p(\theta \mid s^t, \mathbf{H}_t)$ exactly, we approximate this posterior with a learned discriminator $q_\phi(\theta \mid s^t, \mathbf{H}_t)$. By using Jensen's Inequality, replacing $p(\theta \mid s^t, \mathbf{H}_t)$ with $q_\phi(\theta \mid s^t, \mathbf{H}_t)$ gives us a variational lower bound $G(\phi, \psi)$ on our objective $F(\phi)$. Hence, we rewrite (30)

$$\begin{aligned} F(\Psi) &\geq \mathbb{H}[\mathbf{A} \mid s^t, \theta; \mathbf{H}_t] + \mathbb{E}_{\theta \sim p(\theta)} [\log q_\phi(\theta \mid s^t; \mathbf{H}_t) \\ &\quad - \log p(\theta \mid \mathbf{H}_t)] = G(\phi, \psi). \end{aligned} \quad (31)$$

Our goal here is to maximize (31) while simultaneously reducing the average regret (that captures the performance shortfall associated with the suboptimal learning target $\hat{\mathcal{E}}$) in (28). To incorporate the latter aspect, we look at the principle of information-directed sampling (IDS) [47]. IDS is an abstract objective for sequential decision-making agents where, at each time period, an agent computes a policy based on the current history \mathbf{H}_t that minimizes the following ratio of average regret to information:

$$J(\pi_\eta, \hat{\mathcal{E}}) = \frac{\mathbb{E}_\pi \left[d(\mathcal{E}, \hat{\mathcal{E}} \mid \mathbf{H}_t) \right]}{\mathbb{I}(\hat{\mathcal{E}}; \mathbf{a}^t, s_{t+1} \mid \mathbf{H}_t)}. \quad (32)$$

IDS is particularly relevant for CSC since it has the dual objective of (a) maximizing both the QoE by narrowing the gap between its performance and (b) that of an expert policy and enhancing the understanding of the environment \mathcal{E} by utilizing the current extracted semantic state and its history. Unlike in [47], we propose to replace the IDS via

$$\begin{aligned} \left[\pi_{\eta}^{\hat{\mathcal{E}}}, \hat{\theta}, \hat{p}(\mathbf{s}^t \mid \mathbf{y}^t) \right] &= \arg \min_{\pi_{\eta}, \theta, p(\mathbf{s}^t \mid \mathbf{y}^t)} \frac{\mathbb{E}_{\pi} \left[d(\mathcal{E}, \hat{\mathcal{E}} \mid \mathbf{H}_t) \right]}{\mathbb{E}_{\mathbf{y}, \mathbf{s}} G(\phi, \psi)} \\ \text{s.t. } p(E_t(\mathbf{s}^t, \hat{\mathbf{s}}^t) < \delta) &\geq 1 - \epsilon, \end{aligned} \quad (33)$$

where the constraint represents the semantic reliability measure defined earlier. For the denominator term, the expectation is over $p(\mathbf{y}^t)p(\mathbf{s}^t \mid \mathbf{y}^t)$, where $p(\mathbf{s}^t \mid \mathbf{y}^t)$ is the semantic decoder distribution. Next, we look at how to solve (33) using MBRL.

2) *Model Based RL as a Bi-level Optimization*: MBRL is formulated as a bi-level optimization in order to capture the interactions between model and policy learning. Further, we look at the optimization problem for the transition model learning and the imitator policy computation.

$$\mathcal{P}_3 : \underbrace{\max_{\pi_{\eta}} J(\pi_{\eta}, \hat{\mathcal{E}})}_{\text{policy-learner}}, \underbrace{\min_{\mathcal{E}} l(\hat{\mathcal{E}}, \mu_{\mathcal{E}}^{\pi_{\eta}})}_{\text{model-learner}} \text{ subject to } c(\hat{\mathcal{E}}) \leq C \quad (34)$$

We use $\mu_{\mathcal{E}}^{\pi_{\eta}} = \frac{1}{T} \sum_{t=0}^T P(\mathbf{s}^t = s, \mathbf{a}^t = a)$ to denote the average state visitation distribution. $c(\hat{\mathcal{E}}) \leq C$ captures the resource constraints at the imitator. The MBRL can be viewed as a two-level optimization process here, and the technique to solve it is described in Appendix G. At the outer level, the objective is to optimize the policy π_{η} to achieve the best possible performance within the learned model. At the inner level, an optimization problem is solved to minimize the prediction error for \mathcal{E} under the induced state distribution of the policy. This is a bi-level optimization because each objective depends on the parameters of both the problems discussed above. The formulation above decomposes MBRL into policy learning and generative model learning components, emphasizing that they are interdependent and must be addressed together for success. The technique to solve \mathcal{P}_3 is discussed in Appendix G and detailed steps are described in Algorithm 2. Next, we look at the NN architecture that we adopt for various learned components in \mathcal{P}_3 .

C. Generative AI Architecture for “Network State Model”

The imitator model considered here is based on the concept of world models proposed in generative AI [27] and is inspired by our cognitive system. Generative AI is especially advantageous for DT-based SC systems for the following reasons. First, it helps to improve the amount of semantic information acquired by the receiver through accurate modeling of the physical world with the help of DT. Second, as the information to be transmitted becomes more complex, such as 2-D images and videos, and eventually 3-D holograms or even higher-dimensional objects in future connected intelligence systems, generative AI can help the receiver generate more information with less semantics transmitted, compared to classical syntactic communication. It comprises three components that work

together to process received semantic information and make decisions based on past experiences. The first component is the *visual sensory component*, which VAE represents. This component is responsible for decoding the received information received by the agent into a respective network state $\hat{\mathbf{s}}^t$. This decoder distribution is represented by $p(\mathbf{s}^t \mid \mathbf{y}^t, \mathbf{H}_t)$. The second component is the *memory component*, represented by mixture-density network (MDN) combined with a recurrent NN (MDN-RNN). It predicts future states based on historical information, allowing the imitator to anticipate and prepare for potential future events. To anticipate the system’s future states, we use a RNN as a predictive model of future \mathbf{s}^{t+1} vectors. However, since many complex environments are stochastic, we train our RNN to output a probability density function $p(\mathbf{s}^{t+1})$ instead of a deterministic prediction of \mathbf{s}^{t+1} . This allows us to capture the uncertainty inherent in the environment and make more informed decisions based on possible future states. In our approach, we approximate $p(\mathbf{s}^{t+1})$ as a mixture of Gaussian distributions. We then train the memory component (MDN-RNN) to output the probability distribution of the next latent vector \mathbf{s}^{t+1} , based on its current and past information. To be more specific, the RNN, which has N_h hidden units, models $P(\mathbf{s}^{t+1} \mid \mathbf{a}^t, \mathbf{s}^t, \mathbf{h}^t)$, where \mathbf{h}^t is the hidden state of the RNN at time step t . We can adjust a temperature parameter T during sampling to control model uncertainty, as in [27]. We have found that adjusting T is useful for training our controller later. The third and final component is the decision-making component, which is the *controller*. It makes decisions based solely on the representations created by the vision and memory components, enabling the agent to take appropriate actions based on past experiences and current situations. At the time t , the imitator takes action $\mathbf{a}^t \in \mathcal{R}^{N_a}$, where N_a is the dimension of the action space. The proposed imitator model has been named the “*network state model*”, which shares similarities with the concept of world models in AI. However, in contrast to traditional world models in AI, the various generative AI components in the network state model (that mimic the physical environment semantics) are optimized to achieve semantic effectiveness (QoE) as close to 1 as possible, which is critical for a DT-based SC system. The training procedure to optimize the generative AI parameters is detailed in Algorithm 2.

Next, we assess the performance shortfall associated with the suboptimal environment (and hence a suboptimal policy compared to expert agent) learned at the imitator.

Theorem 2. *(Global performance of equilibrium pair)* Consider a pair of policy and environment model at the imitator, $(\pi_{\eta}, \mathcal{E})$, such that simultaneously

$$D_{TV}(P(\cdot \mid \mathbf{s}, \mathbf{a}, \hat{\theta}), \hat{P}(\cdot \mid \mathbf{s}, \mathbf{a}, \hat{\theta})) \leq \epsilon_{\mathcal{E}}, \forall \mathbf{s}, \mathbf{a}, \hat{\theta} \quad (35)$$

and for simplicity, we assume that the extrinsic reward is always zero and the semantic effectiveness (intrinsic reward) is bounded, such that $\mathcal{R}(\mathbf{s}) \leq R_{\max}, \forall \mathbf{s} \in \mathcal{S}$. Also,

Algorithm 2 Proposed Bi-level Optimization for Receiver Design

Given: y
Initialize: Sample $\theta \sim p(\theta)$, and $s^t \sim p(s^t)$. Collect 10,000 rollouts from a random policy

- 1: **for** $(t = 0, 1, 2, \dots)$ **do**
 - 2: Train VAE to decode y^t into $\hat{s}^t \in \mathcal{R}^N$
 - 3: Update the history of observations with \hat{s}^t .
 - 4: Train MDN-RNN to model $p(\theta | \mathbf{H}_t, \hat{s}^t)$. Sample the point estimate $\hat{\theta} \sim \arg \max p(\theta | \mathbf{H}_t, \hat{s}^t)$ as its mode.
 - 5: Train the transition model $p(\hat{s}^{t+1} | \hat{s}^t, \hat{\theta}, \mathbf{H}_t)$ using RNN.
 - 6: Evolve controller to maximize the expected cumulative reward.
 - 7: **end for**
-

assume that the information learned about the environment at the imitator is incremental over time, i.e. $\mathbb{I}(\hat{\mathcal{E}}; \mathbf{a}^t, s^{t+1} | \mathbf{H}_t) = \mathbb{I}_0 t$, where \mathbb{I}_0 is information learned during the first communication instance. Then, we show that

$$|J(\pi_\eta, \mathcal{E}) - J(\pi_\eta, \hat{\mathcal{E}})| \leq \frac{2\gamma\epsilon\mathcal{E}R_{\max}}{(1-\gamma)^2\mathbb{I}_0 t} \quad \forall \pi_\eta \quad (36)$$

Proof: See Appendix E. ■

Theorem 2 indicates that, with time, the discrepancy between the IDS objective and a hypothetical scenario where the imitator has perfect knowledge of \mathcal{E} decreases to zero. This means that as communication between the expert and the imitator continues, the imitator can accurately replicate the expert's policies.

Theorem 3. Suppose we have policy-model pair (π_η, \mathcal{E}) such that the following conditions hold simultaneously:

$$l(\hat{\mathcal{E}}, \mu_{\mathcal{E}}^{\pi, t}) \leq \epsilon_{\mathcal{E}} \quad \forall t \text{ and } J(\pi, \hat{\mathcal{E}}) \geq \sup_{\pi'} J(\pi', \hat{\mathcal{E}}) - \epsilon_{\pi}. \quad (37)$$

Let π^* be an optimal policy so that $J(\pi^*, \mathcal{E}) \geq J(\pi', \mathcal{E}) \quad \forall \pi'$. Then, at time t , the performance of the imitator (evaluated using QoE) will be bounded, as follows:

$$\begin{aligned} J(\pi^*, \mathcal{E}) - J(\pi, \mathcal{E}) &\leq \frac{2R_{\max}}{(1-\gamma)\mathbb{I}_0 t} D_{TV}(\mu_{\mathcal{E}}^{\pi^*}, \mu_{\hat{\mathcal{E}}}^{\pi^*}) \\ &\quad + \frac{2\gamma\sqrt{\epsilon_{\mathcal{E}}}R_{\max}}{(1-\gamma)^2\mathbb{I}_0 t}. \end{aligned} \quad (38)$$

Proof: See Appendix F. ■

Theorem 3 shows that, following a finite number of communication instances, if $\mathbb{I}_0 t$ is a finite value, the discrepancy in transition probability modeling (denoted by $D_{TV}(\mu_{\mathcal{E}}^{\pi^*}, \mu_{\hat{\mathcal{E}}}^{\pi^*})$) dominates the error in the IDS objective with regards to the expert and imitator nodes. Therefore, accurately modeling the physical environment, specifically the network state transitions in our DT-based system, is crucial to ensure that the imitator's policy closely matches the optimal policy generated by the expert. As we observed in Theorem 2, using a generative AI architecture at the imitator can help to close the gap between the physical models of the expert and imitator.

V. SIMULATION RESULTS AND ANALYSIS

To validate the effectiveness of our CSC approach, we must first confirm that our learning methods can accurately identify causal relationships in the world model. We must also ensure that our objective functions and semantic metrics facilitate this process and speed up the discovery of causal relationships. We begin with a dataset $\mathcal{X} = \{\mathbf{x}_s\}_{s=1}^S$ of S samples where each sample \mathbf{x}_s consists of N stationary time-series $\mathbf{x}_s = \{\mathbf{x}_{s,1}, \mathbf{x}_{s,2}, \dots, \mathbf{x}_{s,N}\}$ across time-steps $t = \{1, \dots, T\}$. This dataset represents the states observed at the expert agent. We denote the t -th time-step of the i -th time-series of \mathbf{x}_s as $\mathbf{x}_{s,i}^t \in \mathcal{R}^D$. We consider an SCM captured by an associated DAG $\mathcal{G}_s^{1:T} = \{\mathcal{V}_s^{1:T}, \mathcal{E}_s^{1:T}\}$, underlying the generative process of each sample. The SCM's endogenous (observed) variables are vertices $v_{s,i}^t \in \mathcal{V}_s^{1:T}$ for each time-series i and each time-step t . Every set of incoming edges to an endogenous variable defines inputs to a deterministic function $g_{s,i}^t$ which determines that variable's value. Using DECI, our learned NN is assumed to follow the generation of \mathbf{x}_s^t as $\mathbf{x}_s^{t+1} = \mathbf{g}_s(\mathbf{x}_s^{\leq t}, \mathcal{G}_s, \theta, \mathbf{a}^t) + \mathbf{v}_s^{t+1}$. \mathbf{x}_s^t is the network state (physical environment) observed by the DT. The edges are defined by ordered pairs of vertices $\mathcal{E}_s^{1:T} = \{(v_{s,i}^t, v_{s,j}^{t'})\}$. Here, we consider $t' = t - 1$, and $\mathbf{g}_{s,i}^t = \mathbf{g}_{s,i}^{t'}$. In other words, our causal graphs and process dynamics are invariant across time for each time series sample and can vary across different samples. Our expert agent NNs must model \mathbf{g}_s via the transition probability and the causal graph \mathcal{G}_s . An encoded version \mathbf{z}^t is extracted using a VAE, which gets communicated to the imitator agent. At the receiver, the objective is to recover the causal graph accurately and regenerate the states. The actions at the receiver belong to a discrete set (of cardinality N_a), with a random chosen probability. The entire state space is divided into multiple nonoverlapping subsets, where for each subset there is a unique action probability, which is a discretized Gaussian distribution centered around the mean $\tilde{\mathbf{a}}$, where the mean is distinct for different subsets. We test the CSC system on three datasets: two fully-observed physics simulations (Kuramoto and Particles) and the Netsim dataset of simulated fMRI data that originally appeared in [48] and was utilized for causal discovery in [44]. Throughout the experiments, we compare the performance of the proposed CSC framework to the two following benchmarks. The first is a classical wireless system that directly transmits the received state information (after encoding using a VAE) without any semantic extraction. The second baseline is an SC system that is causality unaware (and hence non-generalizable) [6]. Herein, the SCEs are identified using our proposed approach and further encoded/decoded using transformer modules as in [6] without any causal structure extraction. Fig. 4a presents a scatter plot of the integrated information in the data, which was calculated using (9) and taking into account two-dimensional states \mathbf{s}_i^t . The brightness of each point in the plot corresponds

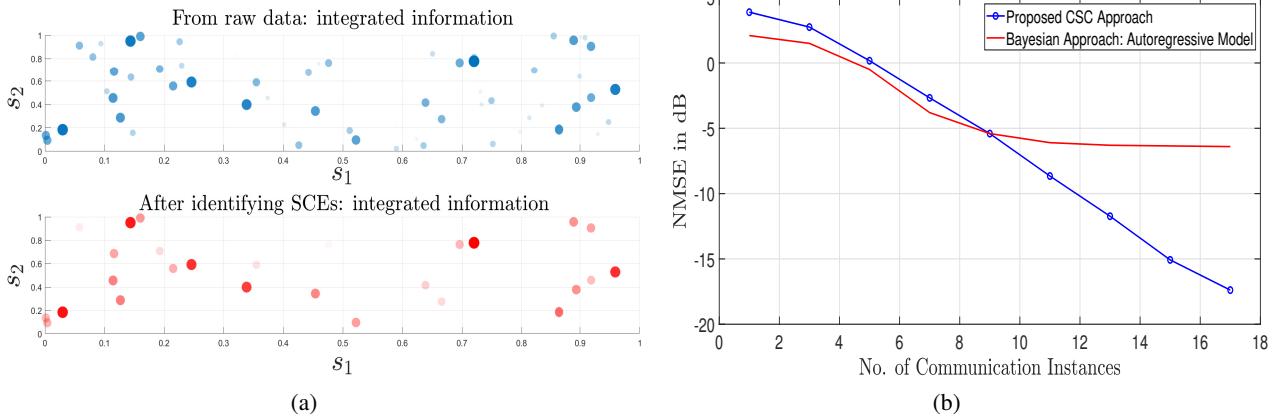


Fig. 4: (a) Integrated information of different subset of states s_i^t . The brightness of the circles represents the magnitude of the integrated information, which is shown in a two-dimensional format. (b) Error in confounding variable estimate between expert and imitator agent.

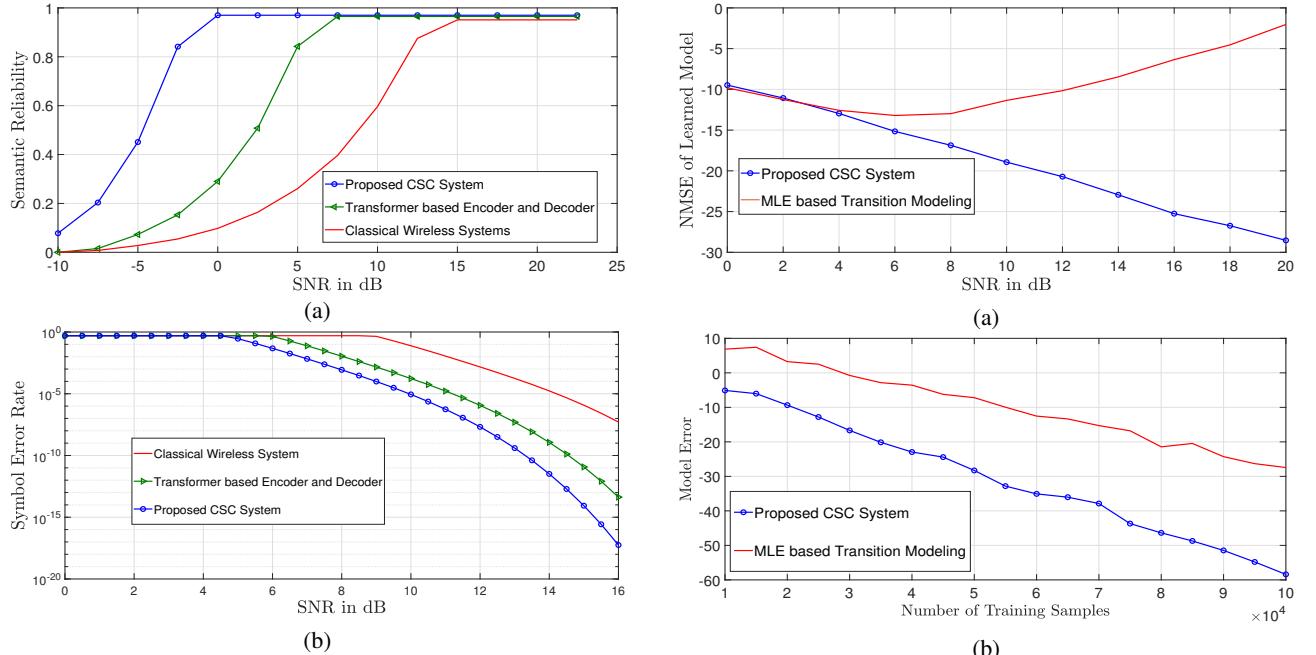


Fig. 5: (a) Semantic reliability vs SNR. (b) Symbol error rate vs SNR.

to the level of integrated information. Additionally, the second row of the same figure shows the SCEs identified through the algorithm described in Section III-C1. The number of SCEs comprises approximately two-thirds of the total number of states present in the data, representing a significant reduction in the amount of data that does convey any semantic information. Fig. 4b illustrates that as communication progresses, the error in confounding variable estimation decreases, in contrast to a model-based Bayesian approach that experiences a performance floor of approximately -5 dB due to inaccurate modeling of the physical environment.

Fig. 5a illustrates the superior performance of the proposed CSC system, which achieves a semantic reliability

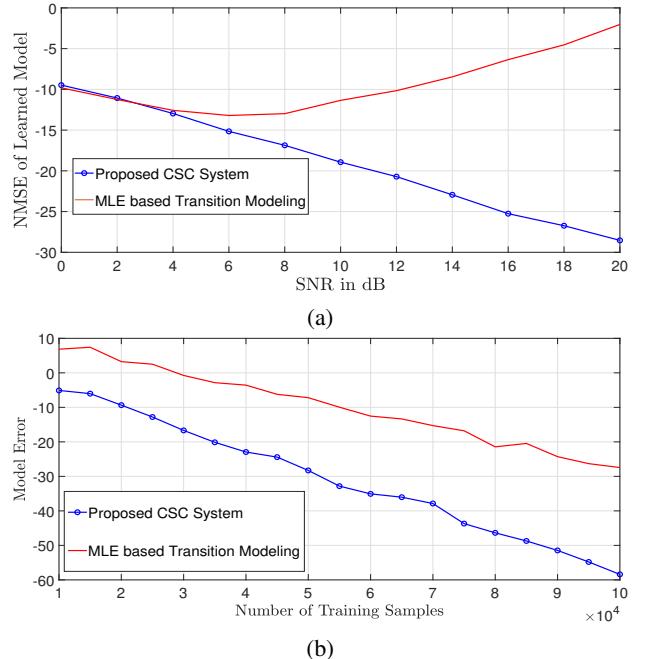


Fig. 6: (a) Environment model error between expert and imitator agents. (b) Generalization performance.

of 0.95 at 0 dB. In contrast, the SC system that does not consider causal structure can achieve the same level of reliability only at 7 dB. This is attributed to the fact that the proposed CSC system requires fewer samples to achieve the desired reliability on the test data set compared to the SC system, which fails to leverage causality. Fig. 5b shows the evaluation results for symbol error rate across different baselines. The results indicate that the proposed CSC system significantly outperforms the traditional SC system, achieving a 10 dB decrease in SER (in dB) at an SNR of 8 dB. Furthermore, compared to classical wireless systems, the proposed CSC system demonstrates a 40 dB decrease in SER, highlighting its superior capabilities in

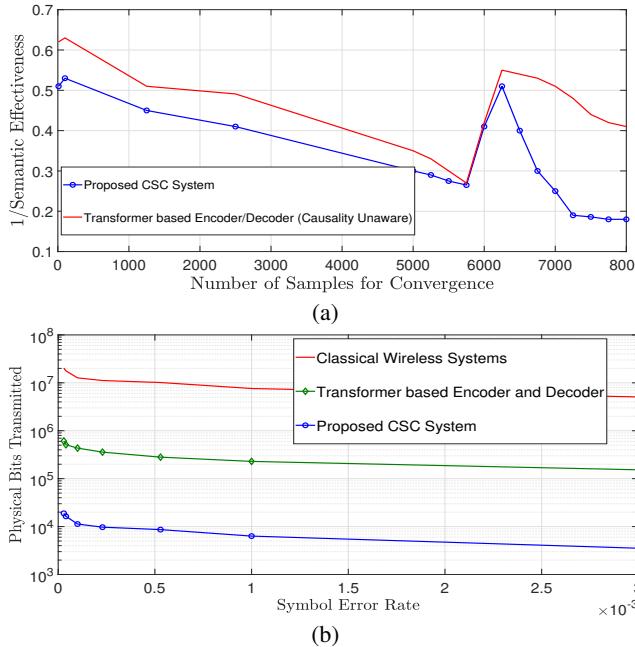


Fig. 7: (a) Semantic effectiveness under non-stationary learning environments (corresponding to change in dynamics g_s). (b) Number of physical bits to be transmitted to achieve a particular SER (corresponding to same semantic reliability).

achieving higher data rates and lower latency.

Fig. 6a presents the evaluation results for normalized mean squared error (NMSE) in the learned model between expert and imitator agents. The results demonstrate that the proposed CSC system significantly outperforms the maximum likelihood (MLE) baselines that utilizes a linear approximation [32] (with just a history of length 1) of the autoregressive model assumed for the dataset, achieving higher model accuracy across all SNRs. This is particularly important for complex datasets, where linear models may perform poorly due to potential nonlinearities. These findings highlight the importance of improving physical model accuracy using advanced AI algorithms, such as causal discovery. In Fig. 6b, we evaluate the generalization performance of the proposed MBRL scheme at the imitator compared to MLE based learner. Here, we computed the model error ($g_s^E - g_s^{\eta}$) between expert and imitator agents for the proposed CSC and the MLE. We observe that our MBRL performs much better than the standard model learner (MLE), especially when the number of samples available is low. This is attributed to the causal structure extraction of the environment dynamics, which leads to better generalizability.

Fig. 7a shows the performance of the proposed CSC system and the causality unaware baseline in a non-stationary learning environment. The goal here is to demonstrate the quick adaptability (with minimal training effort compared to state-of-the-art) of the CSC approach proposed to changes in the transition dynamics g_s of the time series data used in the simulations. In this case, the transition dynamics g_s change after around 5.7K samples. The results demon-

strate that the proposed CSC system efficiently recovers from dynamics perturbations, requiring approximately three times fewer samples compared to the causality-unaware system, which necessitates more retraining efforts. Hence, the proposed approach is effective for generalizing to multiple environments. In Fig. 7b, we observe that the CSC system outperforms conventional communication systems and other SC baselines in terms of efficiency (measured in the number of bits required to convey the same amount of semantic information). Specifically, we analyze the number of bits transmitted to achieve a particular SER for 10,000 communication instances. The CSC system requires significantly fewer bits (by a factor of 1000) compared to the classical system (without reasoning part), highlighting the significance of our proposed approach. Furthermore, when compared to state-of-the-art algorithms that do not consider causal reasoning, the proposed CSC approach is more robust to SER.

VI. CONCLUSION

In this paper, for the first time in the literature, we have presented a new vision of a DT-based SC system entitled CSC that relies upon the recently emerged theory of consciousness measures based on IIT and AI tools such as IL and MBRL. This approach enables causal discovery of the network state, allowing it to be generalized across multiple wireless environments. We have formulated a bi-level optimization based on variational inference and information bottleneck principle to learn causal discovery, state transitions, and semantic representation at the transmitter (expert agent). The optimized semantic representation is such that the receiver is able to extract maximum semantic information while at the same time revealing minimum information about the state. At the receiver (the imitator node in IL), using a generative AI architecture originally proposed for “world models” in RL, the node improves its knowledge about the network state transitions and causality over time. The proposed bi-level optimization is formulated using the principles of variational inference and IDS. We have shown analytically the performance shortfall associated with a suboptimal environment learning at the imitator. Simulation results demonstrate our proposed CSC’s superiority in improving communication efficiency (minimal transmission) and reliability compared to classical communication and state-of-the-art SC systems.

REFERENCES

- [1] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, “Digital Twin in Industry: State-of-the-Art,” *IEEE Transactions on Industrial Informatics*, vol. 15, no. 4, pp. 2405–2415, Apr. 2019.
- [2] L. U. Khan, Z. Han, W. Saad, E. Hossain, M. Guizani, and C. Seon Hong, “Digital Twin of Wireless Systems: Overview, Taxonomy, Challenges, and Opportunities,” *IEEE Communications Surveys and Tutorials*, vol. 24, no. 4, 2022.
- [3] M. Kountouris and N. Pappas, “Semantics-empowered communication for networked intelligent systems,” *IEEE Communications Magazine*, vol. 59, no. 6, pp. 96–102, Jun. 2022.

- [4] E. Uysal, O. Kaya, A. Ephremides, J. Gross, M. Codreanu, P. Popovski, M. Assad, G. Liva, A. Munari, B. Soret, and T. Soleymani, "Semantic Communications in Networked Systems: A Data Significance Perspective," *IEEE Network*, vol. 36, no. 4, pp. 233–240, Oct. 2022.
- [5] C. Chaccour, W. Saad, M. Debbah, Z. Han, and H. V. Poor, "Less Data, More Knowledge: Building Next Generation Semantic Communication Networks," *arXiv preprint arXiv:2211.14343*, Nov. 2022.
- [6] H. Xie, Z. Qin, G. Y. Li, and B-H. Juang, "Deep Learning Enabled Semantic Communication Systems," *IEEE Transactions on Signal Processing*, vol. 69, pp. 2663–2675, Apr. 2021.
- [7] S. Seo, J. Park, S-W Ko, J. Choi, M. Bennis, and S-L Kim, "Towards Semantic Communication Protocols: A Probabilistic Logic Perspective," *arXiv:2207.03920v1*, Jul. 2022.
- [8] J. Liu, W. Zhang, and H. V. Poor, "A Rate-Distortion Framework for Characterizing Semantic Information," in *Proceedings of IEEE International Symposium on Information Theory (ISIT)*, Grenoble, France, Jul. 2021.
- [9] M. K. Farshbafan, W. Saad, and M. Debbah, "Curriculum Learning for Goal-Oriented Semantic Communications with a Common Language," *arXiv preprint arXiv:2111.08051*, Feb. 2022.
- [10] M. K. Farshbafan, W. Saad, and M. Debbah, "Common Language for Goal-Oriented Semantic Communications: A Curriculum Learning Framework," in *Proceedings of IEEE International Conference on Communications (ICC)*, May 2022.
- [11] Y. E. Sagduyu, T. Erpek, S. Ulukus, and A. Yener, "Is Semantic Communications Secure? A Tale of Multi-Domain Adversarial Attacks," *arXiv preprint arXiv:2212.10438*, Dec. 2022.
- [12] W. Sun, S. Lei, L. Wang, Z. Liu, and Y. Zhang, "Adaptive Federated Learning and Digital Twin for Industrial Internet of Things," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 8, pp. 5605–5614, Aug. 2021.
- [13] Y. Lu, X. Huang, K. Zhang, S. Maharjan, and Y. Zhang, "Low-latency Federated Learning and Blockchain for Edge Association in Digital Twin Empowered 6G Networks," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 7, pp. 5098–5107, Jul. 2021.
- [14] O. Hashash, C. Chaccour, W. Saad, K. Sakaguchi, and T. Yu, "Towards a decentralized metaverse: Synchronized orchestration of digital twins and sub-metaverses," in *IEEE International Conference on Communications (ICC), Next-generation Networking and Internet Symposium*, Rome, Italy, May 2023.
- [15] Y. Dai, K. Zhang, S. Maharjan, and Y. Zhang, "Deep Reinforcement Learning for Stochastic Computation Offloading in Digital Twin Networks," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 7, pp. 4968–4977, Jul. 2021.
- [16] Y. Lu, S. Maharjan, and Y. Zhang, "Adaptive Edge Association for Wireless Digital Twin Networks in 6G," *IEEE Internet of Things Journal*, vol. 8, no. 22, pp. 16219–16230, Nov. 2021.
- [17] D. V. Huynh, S. R. Khosravirad, A. Masaracchia, O. A. Dobre, and T. Q. Duong, "Edge Intelligence-Based Ultra-Reliable and Low-Latency Communications for Digital Twin-Enabled Metaverse," *IEEE Wireless Communications Letters*, vol. 11, no. 8, Aug. 2022.
- [18] C. Ruah, O. Simeone, and B. Al-Hashimi, "A Bayesian Framework for Digital Twin-Based Control, Monitoring, and Data Collection in Wireless Systems," *arXiv preprint arXiv:2212.01351*, 2023.
- [19] O. Takayuki, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters, "An Algorithmic Perspective on Imitation Learning," *Foundations and Trends in Robotics*, vol. 7, no. 1-2, pp. 1–179, Mar. 2018.
- [20] Y. Xiao, Z. Sun, G. Shi, and D. Niyato, "Imitation Learning-based Implicit Semantic-aware Communication Networks: Multi-layer Representation and Collaborative Reasoning," *IEEE Journal on Selected Areas in Communications*, vol. 43, no. 1, Mar. 2023.
- [21] R. Vuorio, J. Brehmer, H. Ackermann, D. Dijkman, T. Cohen, and P. de Haan, "Deconfounded Imitation Learning," *arXiv preprint arXiv:2211.02667*, 2022.
- [22] J. Pearl, "Causal inference," *International Journal on Biostat*, vol. 2, pp. 39–58, 2010.
- [23] L. Albantakis, L. Barbosa, G. Findlay, M. Grasso, A. M. Haun, W. Marshall, Mayner W. G, A. Zaeemzadeh, M. Boly, and B. E. Juel and S. Sasai, "Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms," *arXiv preprint arXiv:2212.14787*, 2022.
- [24] A. Haun and G. Tononi, "Why does space feel the way it does? Towards a principled account of spatial experience," *Entropy*, vol. 21, no. 12, 2019.
- [25] Y. Matsuo, Y. LeCun, M. Sahani, D. Precup, D. Silver, M. Sugiyama, E. Uchibe, M. Sugiyama, E. Uchibe, and J. Morimoto, "Deep Learning, Reinforcement Learning, and World Models," *Neural Networks, Special issue on Artificial Intelligence and Brain Science*, vol. 152, pp. 267–275, Aug. 2022.
- [26] V. D. Nguyen, T. X. Vu, N. T. Nguyen, D. C. Nguyen, M. Juntti, N. C. Luong, D. T. Hoang, D. N. Nguyen, and S. Chatzinotas, "Network-Aided Intelligent Traffic Steering in 6G ORAN: A Multi-Layer Optimization Framework," *arXiv preprint arXiv:2302.02711*, Feb. 2023.
- [27] D. Ha and J. Schmidhuber, "Recurrent World Models Facilitate Policy Evolution," in *Advances in Neural Information Processing Systems 31 (NeurIPS)*, 2018.
- [28] T. Geffner, J. Antoran, A. Foster, W. Gong, C. Mao, E. Kiciman, A. Sharma, A. Lamb, M. Kukla, M. Pawlowski, M. Allamanis, and C. Zhang, "Deep End-to-end Causal Inference," *arXiv preprint arXiv:2202.02195*, 2022.
- [29] T. M. Moerland, J. Broekens, A. Plaat, and C. M. Jonker, "Model-based reinforcement learning: A survey," *Foundations and Trends® in Machine Learning*, vol. 16, no. 1, pp. 1–118, 2023.
- [30] Pedro A Ortega, Markus Kunesch, Grégoire Deléring, Tim Graweinein, Jordi Grau-Moya, Joel Veness, Jonas Buchli, Jonas Degraeve, Bilal Piot, Julien Perolat, Tom Everitt, Corentin Tallec, Emilio Parisotto, Tom Erez, Yutian Chen, Scott Reed, Marcus Hutter, Nando de Freitas, and Shane Legg, "Shaking the foundations: delusions in sequence models for interaction and control," *arXiv preprint arXiv:2110.10819*, 2021.
- [31] B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio, "Toward Causal Representation Learning," *Proceedings of the IEEE*, vol. 109, no. 5, Feb. 2021.
- [32] M. Tomar, A. Zhang, R. Calendra, M. E. Taylor, and J. Pineau, "Model-Invariant State Abstractions for Model-Based Reinforcement Learning," *arXiv preprint arXiv:2102.09850*, 2021.
- [33] C. K. Thomas and W. Saad, "Neuro-Symbolic Causal Reasoning Meets Signaling Game for Emergent Semantic Communications," *IEEE Transactions on Wireless Communications*, Oct. 2023.
- [34] D. Pathak, P. Agrawal, A. Efros, and T. Darrell, "Curiosity-driven Exploration by Self-supervised Prediction," in *In International conference on machine learning (ICML), PMLR*, 2017, pp. 2778–2787.
- [35] M. Oizumi, L. Albantakis, and G. Tononi, "From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0," *PLoS computational biology*, vol. 10, no. 5, 2014.
- [36] D. Abel, D. Arumugam, L. Lehnert, and M. Littman, "State Abstractions for Lifelong Reinforcement Learning," in *International Conference on Machine Learning, PMLR*, July 2018, pp. 10–10.
- [37] I. Higgins et al., "β-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework," in *ICLR*, Apr. 2017.
- [38] Y. Shao, Q. Cao, and D. Gunduz, "A Theory of Semantic Communication," *arXiv preprint arXiv:2212.01485*, 2022.
- [39] A. Hatcher, "Algebraic Topology," in *Cambridge University Press: Cambridge, UK*, 2002.
- [40] O. Caramello, "Theories, Sites, Toposes: Relating and Studying Mathematical Theories Through Topos-theoretic 'bridges,'" *Oxford University Press*, 2018.
- [41] X. Zheng, B. Aragam, P. Ravikumar, and E. P. Xing, "DAGs with no tears: Continuous optimization for structure learning," *arXiv preprint arXiv:1803.01422*, 2018.
- [42] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational Inference: A Review for Statisticians," *Journal of the American statistical Association*, vol. 112, no. 518, pp. 859–877, 2017.
- [43] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The Graph Neural Network Model," *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 561–80, Dec. 2008.
- [44] S. Löwe, D. Madras, R. Zemel, and M. Welling, "Amortized Causal Discovery: Learning to Infer Causal Graphs from Time-Series Data," *Proceedings of Machine Learning Research*, vol. 140, pp. 509–525, Jun. 2022.

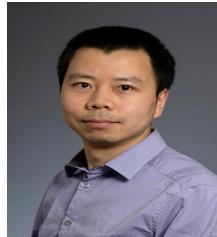
- [45] E. Beck, C. Bockelmann, and A. Dekorsy, “Semantic Information Recovery in Wireless Networks,” *arXiv:2204.13366v3*, Mar. 2023.
- [46] A. Rajeswaran, K. Lowrey, E. Todorov, and S. Kakade, “Towards Generalization and Simplicity in Continuous Control,” in *Neural Information Processing Systems*, 2017.
- [47] D. Arumugam and B. V. Roy, “The value of information when deciding what to learn,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 9816–9827, May 2021.
- [48] T. Kipf, E. Fetaya, K-C. Wang, M. Welling, and R. Zemel, “Neural Relational Inference for Interacting Systems,” *arXiv preprint arXiv:1802.04687*, 2018.



Walid Saad (S’07, M’10, SM’15, F’19) received his Ph.D degree from the University of Oslo, Norway in 2010. He is currently a Professor at the Department of Electrical and Computer Engineering at Virginia Tech, where he leads the Network sciENCE, Wireless, and Security (NEWS) laboratory. He is also the Next-G Wireless Faculty Lead at Virginia Tech’s Innovation Campus. His research interests include wireless networks (5G/6G/beyond), machine learning, game theory, security, UAVs, semantic communications, cyber-physical systems, and network science. Dr. Saad is a Fellow of the IEEE. He is also the recipient of the NSF CAREER award in 2013, the AFOSR summer faculty fellowship in 2014, and the Young Investigator Award from the Office of Naval Research (ONR) in 2015. He was the (co-)author of eleven conference best paper awards at IEEE WiOpt in 2009, ICIMP in 2010, IEEE WCNC in 2012, IEEE PIMRC in 2015, IEEE SmartGridComm in 2015, EuCNC in 2017, IEEE GLOBECOM (2018 and 2020), IFIP NTMS in 2019, IEEE ICC (2020 and 2022). He is the recipient of the 2015 and 2022 Fred W. Ellersick Prize from the IEEE Communications Society, of the IEEE Communications Society Marconi Prize Award in 2023, and of the IEEE Communications Society Award for Advances in Communication in 2023. He was also a co-author of the papers that received the IEEE Communications Society Young Author Best Paper award in 2019, 2021, and 2023. Other recognitions include the 2017 IEEE ComSoc Best Young Professional in Academia award, the 2018 IEEE ComSoc Radio Communications Committee Early Achievement Award, and the 2019 IEEE ComSoc Communication Theory Technical Committee Early Achievement Award. From 2015-2017, Dr. Saad was named the Stephen O. Lane Junior Faculty Fellow at Virginia Tech and, in 2017, he was named College of Engineering Faculty Fellow. He received the Dean’s award for Research Excellence from Virginia Tech in 2019. He was also an IEEE Distinguished Lecturer in 2019-2020. He has been annually listed in the Clarivate Web of Science Highly Cited Researcher List since 2019. He currently serves as an Area Editor for the IEEE Transactions on Network Science and Engineering and the IEEE Transactions on Communications. He is the Editor-in-Chief for the IEEE Transactions on Machine Learning in Communications and Networking.



Christo Kurisummoottil Thomas (S’17, M’20) received the B.S. degree in electronics and communication engineering from the National Institute of Technology, Calicut, India, in 2010, the M.S. degree Institute of Science, Bengaluru, India, in 2012, and the Ph.D. degree from EUR-ECON, France, in 2020. From 2012 to 2014, he was a Staff Design Engineer on 4G LTE at Broadcom Communications, Bengaluru. From 2014 to 2017, he was a Design Engineer at Intel Corporation, Bengaluru. From November 2020 to June 2022, he was a Staff Engineer on 5G modems at the Wireless Research and Development Division, Qualcomm Inc., Espoo, Finland. He is currently a Post-Doctoral Fellow with the Department of Electrical and Computer Engineering, Virginia Tech. His research interests include semantic communications, statistical signal processing, and machine learning for wireless communications. He was a recipient of the Best Student Paper Award from the IEEE SPAWC 2018, Kalamata, Greece. He also received third prize for his team titled “Learned Chester” ML5G-PHY channel estimation challenge, as part of the ITU AI/ML in 5G challenge, conducted at NCSU, USA, in 2020.



Yong Xiao (S’06, M’12) received his B.S. degree in electrical engineering from China University of Geosciences, Wuhan, China in 2002, M.Sc. degree in telecommunication from Hong Kong University of Science and Technology in 2006, and his Ph. D degree in electrical and electronic engineering from Nanyang Technological University, Singapore in 2012. He is now a professor in the School of Electronic Information and Communications at the Huazhong University of Science and Technology (HUST), Wuhan, China. He is also with Peng Cheng Laboratory, Shenzhen, China and Pazhou Laboratory (Huangpu), Guangzhou, China. He is the associate group leader of the network intelligence group of IMT-2030 (6G promoting group) and the vice director of 5G Verticals Innovation Laboratory at HUST. Before he joins HUST, he was a research assistant professor in the Department of Electrical and Computer Engineering at the University of Arizona where he was also the center manager of the Broadband Wireless Access and Applications Center (BWAC), an NSF Industry/University Cooperative Research Center (I/UCRC) led by the University of Arizona. His research interests include machine learning, game theory, distributed optimization, and their applications in semantic communications, semantic-aware networking, cloud/fog/mobile edge computing, green communication systems, and Internet-of-Things (IoT).